

MAGAZINE

BSD

FOR NOVICE AND ADVANCED USERS

**IMPLEMENTING IN-MEMORY CACHE
IN THE BEAST ARCHITECTURE**

DOCKER CLEANUP

INTERVIEW WITH ALEXANDAR S. SOKOLOVSKI

FOUNDER OF SCI BI

**FREENAS GETTING STARTED GUIDE
PART 2 THE INITIAL CONFIGURATION WIZARD**

**MySQL SERVER 5.7
STARTUP PROBLEM ON FREEBSD 10.2**

**LET US START AT THE VERY BEGINNING
WITH OPEN SOURCE**

**My Top 9 Favorite
Python Deep Learning Libraries**

VOL 10 NO 06

ISSUE 06/2016 (82)

1898-9144

THE FREEENAS MINI XL HAS ARRIVED



ENTERPRISE-CLASS HARDWARE, RUNNING
THE WORLD'S MOST POPULAR OPEN SOURCE
STORAGE OPERATING SYSTEM.

For more information on the FreeNAS Mini,
visit ixsystems.com/mini today.

EDITORS' WORD

Dear Readers,

We hope you have been doing well. Holiday season has started, both public holidays and vacation. How was the 4th of July in the USA and Canada's day? 1st of July? Did you have a great time with your friends and families? If you are not based in North America, let us know what and when is the most important holiday in your country!

We all love holidays, but let's dive into this month's issue. As always, we will start with news from the industry. There has been a lot going on in the open source world.

Mikhail Zakharov shared the third article from his storage series. This time you will read "Implementing In-Memory Cache in the BeaST Architecture". It's the next step in building a perfect, dual-headed FreeBSD storage concept, but no worries, there will be more in an upcoming issue.

If you have ever had a problem with MySQL Server Startup on FreeBSD, Andrey Ferriyan can help you out with his article.

Our expert, E.G. Nadhan from Red Hat wrote a great article about the open source community. Let us know your thoughts about your favorite benefit from using open source. Nicolas de Loof shared how to do a "Docker Cleanup" and Adrian Rosebrock his "Top (not) 9 Favorite Python Deep Learning Libraries" with us.

If you have always wanted to start your journey with FreeNAS, take a look at Mark Von Fange's second article about "FreeNAS Getting Started Guide: Part 2, The Initial Configuration Wizard". Pure knowledge from the FreeNAS source.

To wrap this issue up, we have a great interview with Aleksandar S. Sokolovski, Founder of SciBi; and, as always, Rob Somerville gets right to the point in his column.

Enjoy your holidays and let us know what topics you would like us to cover in upcoming issues.

Marta & BSD Team



Editor in Chief:

Marta Ziemianowicz

marta.ziemianowicz@software.com.pl

Contributing:

Mikhail E. Zakharov, Andrey Ferriyan, E.G. Nadhan, Nicolas De Loof, Adrian Rosebrock, Mark VonFange, Aleksandar S. Sokolovski, and Rob Somerville.

Top Beta testers & Proofreaders:

Annie Zhang, Denise Ebery, Eric Geissinger, Luca Ferrari, Imad Soltani, Olaoluwa Omokanwaye, Radjis Mahangoe, Mani Kanth and Mark VonFange.

Special Thanks:

Annie Zhang

Denise Ebery

DTP:

Marta Ziemianowicz

Senior Consultant/Publisher:

Pawel Marciniaik

pawel@software.com.pl

CEO:

Joanna Kretowicz

joanna.kretowicz@software.com.pl

Publisher:

Hakin9 Media SK 02-676 Warsaw, Poland Postepu 17D Poland worldwide publishing editors@bsdmag.org www.bsdmag.org

Hakin9 Media SK is looking for partners from all over the world. If you are interested in cooperation with us, please contact us via e-mail: editors@bsdmag.org.

All trademarks presented in the magazine were used only for informative purposes. All rights to trademarks presented in the magazine are reserved by the companies which own them.

CONTENTS

News

BSD World Monthly News **5**

by *Marta Ziemianowicz*

This column presents the latest news coverage of events, product releases and trending topics.

FreeBSD Corner

Implementing In-Memory Cache in the BeaST Architecture **18**

by *Mikhail E. Zakharov*

Recently we have created the reliable, dual-headed FreeBSD based storage concept. Then we improved it with the famous ZFS feature. Now we will try to implement a shared in-memory cache to make our model even more serious.

MySQL Server 5.7 Startup Problem on FreeBSD 10.2 **35**

by *Andrey Ferriyan*

I've just installed MySQL server 5.7 for FreeBSD 10.2 using pkg package manager. It seems the problem occurred when I went to start MySQL for the first time. This time I'm using the real VPS server and not VirtualBox.

Expert Says

Let Us Start at the Very Beginning with Open Source **40**

by *E.G. Nadhan*

The Red Hat coloring booth at SX Create during the South by Southwest (SXSW) conference and festival helped to tell the open source story of collaboration in a way that was simple enough for children to understand. It was a booth where new ideas manifested themselves in the form of colorful shapes and forms on the walls. Thousands of attendees contributed, bringing the 12 foot tall space to life. Ideas were expanded on by other children, modified or cloned by some to give birth to other related ideas. Continuous, enthusiastic collaboration at its peak in an environ-

ment ripe for open communication -- and thus, innovation. Sound familiar?

Docker

Docker Cleanup **43**

by *Nicolas De Loof*

Most Docker newcomers are disappointed when, at some time, after various experiments with Docker, they hit a no space left on device issue. Docker does indeed store all containers, images and volumes in /var/lib/docker, which can quickly grow to gigabytes.

Python

My Top 9 Favorite Python Deep Learning Libraries **46**

by *Adrian Rosebrock*

So you're interested in deep learning and Convolutional Neural Networks. But where do you start? Which library do you use? There are just so many!

This list is by no means exhaustive, it's simply a list of libraries that I've used in my computer vision career and found particularly useful at one time or another.

FreeNAS

FreeNAS Getting Started Guide: Part 2, The Initial Configuration Wizard **57**

by *Mark VonFange*

This article series is intended to serve as an introductory guide to assist FreeNAS users in planning, installation, configuration and administration for their FreeNAS storage systems. This month's article will cover getting your FreeNAS set up using the Initial Configuration Wizard.

Interview

Interview with Aleksandar S. Sokolovski, Co-Founder of Sci Bi **69**

by *Marta Ziemianowicz, Marta Strzelec & Marta Siennicka*

CONTENTS

Rob's Column **78**

by Rob Somerville

With the recent British referendum decision to exit from the Europe Union, one of the few opinion polls to correctly predict the outcome was online. What impact is the Internet having on the established political order and is this a force for good?

BSD Certification

The BSD Certification Group Inc. (BSDCG) is a non-profit organization committed to creating and maintaining a global certification standard for system administration on BSD based operating systems.

WHAT CERTIFICATIONS ARE AVAILABLE?

BSDA: Entry-level certification suited for candidates with a general Unix background and at least six months of experience with BSD systems.

BSDP: Advanced certification for senior system administrators with at least three years of experience on BSD systems. Successful BSDP candidates are able to demonstrate strong to expert skills in BSD Unix system administration.

WHERE CAN I GET CERTIFIED?

We're pleased to announce that after 7 months of negotiations and the work required to make the exam available in a computer based format, that the BSDA exam is now available at several hundred testing centers around the world. Paper based BSDA exams cost \$75 USD. Computer based BSDA exams cost \$150 USD. The price of the BSDP exams are yet to be determined.

Payments are made through our registration website:
<https://register.bsdcertification.org//register/payment>

WHERE CAN I GET MORE INFORMATION?

More information and links to our mailing lists, LinkedIn groups, and Facebook group are available at our website:
<http://www.bsdcertification.org>

Registration for upcoming exam events is available at our registration website:
<https://register.bsdcertification.org//register/get-a-bsdcg-id>

PC-BSD's Lumina Desktop Now In Beta For v1.0

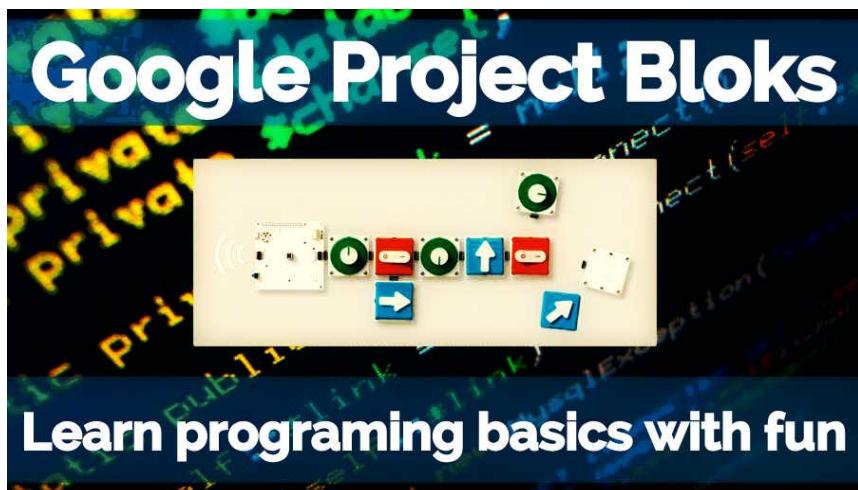
The Lumina Desktop Environment has made available their v1.0 beta release of the Qt-written desktop.

PC-BSD developers and others continue working on Lumina as an alternative, light-weight desktop environment. While originating in the BSD world, Lumina continues to be designed to work on any Unix-like OS and is licensed under a 3-clause BSD license. Should you not be familiar with Lumina from our past articles, visit Lumina-Desktop.org to learn more about the project.

Lumina 1.0-Beta 1 was released to start off July. The release announcement for this version simply reads, "This is a beta-quality release of the Lumina Desktop Environment (1.0.0-Devel, snapshot on July 1 2016). Although highly stable and ready for use on most desktop systems, many backend systems are still changing which may cause the user's settings to be reset between versions (although this is minimized as much as possible)."

http://www.phoronix.com/scan.php?page=news_item&px=Lumina-1.0-Beta

Google Launches ‘Project Bloks’ To Help Young Learners Learn To Code



Based on the principles of tangible programming, Google has devised a new learning tool called ‘Project Bloks’. With the help of three components, Bloks creates a ‘physical’ program that teaches the coding basics to young learners.

Google calls programming a way to develop creativity and computational thinking. This view is supported by different governments from all around the world — some of them have even included coding as

a compulsory subject in school curriculums.

In its latest efforts to make coding more accessible and easier, Google has partnered with the Stanford University and design firm IDEO to start a new initiative called ‘Project Bloks’.

It’s an open platform created to help the kids learn programming. With the help of physical components, kids use ‘Bloks’ as a reference design to create a physical ‘program’ that controls real world objects.

Specifically targeted at kids, this system consists of 3 core blocks — Brain Board, Base Board, and Pucks.

The Raspberry Pi-based Brain Board acts as the main controller that also houses power and communication parts.

The Base Boards is a small block that plays the role of a home that provides LED, haptic, and audio feedback.

The little Pucks are control modules that are made to perform a single function, like on and off, increasing/decreasing volume, stop, go, etc. When connected to a base board or another case board, a puck sends the specific commands to the brain board.

With Project Bloks, Google intends to teach kids the logic behind coding and simple fundamentals. With this physical coding initiative, kids can grasp the basic skills and transfer them to the real-world applications.

Let’s hope that Google makes coding toys cheaper and develops Bloks into a powerful toy that developers can use to create customizable learning tools.

<http://fossbytes.com/google-project-bloks-coding/>

Mozilla Pushes Online Privacy with New Open Source Funding Awards

Mozilla is funneling yet more money into the open source ecosystem. This week, the organization, best known for the Firefox Web browser, announced an award of \$385,000 to fund eight open source projects, including several important online privacy platforms.

The awards are part of the Mozilla Open Source Support (MOSS) program. Alongside initiatives like the funding of gigabit Internet, MOSS is an example of how big-name open source foundations are pursuing new strategies for shaping the open source ecosystem with their cash reserves.

The awards announced this week will support the following projects:

- Tor, which develops software for hiding your identity and location online.
- Tails, a Linux-based operating system designed to provide out-of-the-box privacy and security.
- Caddy, a Web server that delivers HTTPS encryption by default.
- Mio, a software programming library.
- DNSSEC/DANE Chain Stapling, which is working to improve the process of secure DNS requests.
- Godot Engine, a game engine.
- PeARS, a private Internet search engine.
- NVDA, a screen reader for Windows.

A number of these projects focus on improving online privacy and security. Mozilla is making pretty clear its commitment to helping to create a more private, decentralized Internet free of censorship and tracking -- or at least one in which censorship and tracking are more difficult than they have become in most online contexts today.

That should send a signal to other open source organizations and developers about the direction in which the money is flowing inside the open source ecosystem. Demand for better privacy and security delivered through open source code is high, and prestigious foundations like Mozilla are putting up the cash to reward developers who deliver those features.

<http://thevar guy.com/open-source-application-software-companies/mozilla-pushes-online-privacy-new-open-source-funding-awa>

ubuntuBSD 16.04 to feature a combo of BusyBox and OpenRC, no systemd



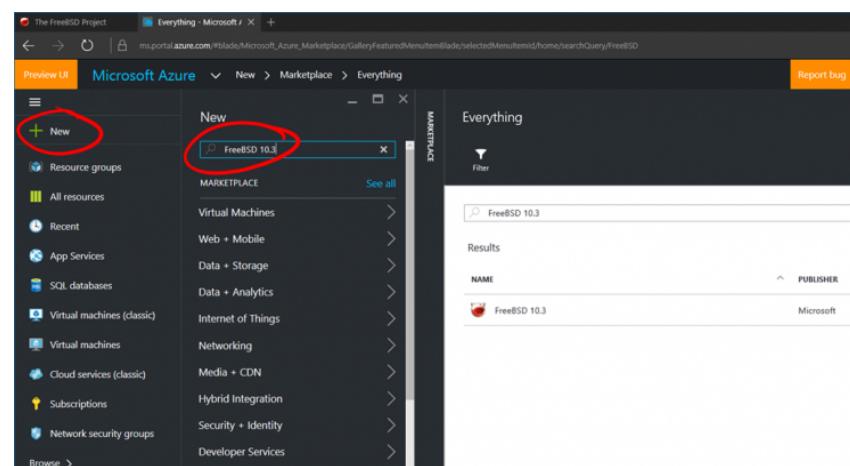
For those interested in the latest developments of ubuntuBSD, the developer mentions that the official release will come bundled with BusyBox and OpenRC, sans systemd.

In a series of tweets, ubuntuBSD project leader Jon Boden has announced a few of the technical features coming to the soon-to-be-released ubuntuBSD 16.04 operating system.

A week ago, we wrote an exclusive story to tell you that the first and major release of the ubuntuBSD OS is coming soon, based on the Ubuntu 16.04 LTS (Xenial Xerus) distribution, but using the powerful kernel from the BSD-based FreeBSD 10.3 operating system.

<https://www.freebsdnews.com/2016/06/24/ubuntubsd-16-04-will-feature-a-combination-of-busybox-and-openrc-but-no-systemd/Upstream>

FreeBSD is available in Microsoft's Azure Marketplace



Microsoft has announced that FreeBSD 10.3 is now available as a virtual machine image in the Azure Marketplace. Azure is Microsoft's cloud computing platform and infrastructure for deploying various applications.

Today, I'm excited to announce the availability of FreeBSD 10.3 as a ready-made VM image available directly from the Azure Marketplace. This means that

not only can you quickly bring-up a FreeBSD VM in Azure, but also that in the event you need technical support, Microsoft support engineers can assist.

Here's how easy it is to get up and going through the Azure portal. Simply click on the '+New' on the left pane (or the marketplace tile on your dashboard), type "FreeBSD 10.3" in the search text box, and you're there.

<https://www.freebsdnews.com/2016/06/17/microsoft-unveils-freebsd-build-for-virtual-appliances-among-azure-updates/>

FreeBSD 11.0 Alpha 4 Released

The fourth alpha release of the upcoming FreeBSD 11.0 is now available for testing.

FreeBSD 11.0 Alpha 4 ships the very latest fixes for this major BSD update. FreeBSD 11.0 is scheduled to be officially released in early September with the code freeze happening last week, the beta builds beginning in July, and release candidates in August. The FreeBSD 11.0 schedule can be found via FreeBSD.org.

FreeBSD 11.0 is bringing updated KMS drivers, Linux binary compatibility layer improvements, UEFI improvements, Bhyve virtualization improvements, and a wide range of other enhancements. FreeBSD 11.0 tentative release notes can be seen here for those wanting to learn more about this forthcoming OS update.

The latest FreeBSD development builds can be found at FreeBSD.org FTP.

http://www.phoronix.com/scan.php?page=news_item&px=FreeBSD-11.0-Alpha-4

Bringing Raspberry Pi to schools in Tanzania



Thanks to open source software, Powering Potential and the Raspberry Pi Foundation are able to bring computers and a library of digital education content to rural schools in the East African nation of Tanzania. Recently, the Foundation funded a project that is now distributing Raspberry Pi computers with uploaded educational content alongside portable projectors and screens to 56 schools across the Zanzibar archipelago and two mainland regions of Tanzania.

The Segal Family Foundation also provided matching funds, which enables the project to reach twice as many schools.

With a five-fold increase in the number of students in the decade following 2003, Tanzania is struggling to provide more schools, classrooms, teachers, desks, and textbooks. Yet whenever you visit rural secondary schools in Tanzania, you will find eager girls and boys in roughly equal numbers outfitted in uniforms with ready smiles.

<https://opensource.com/education/16/6/interview-janice-lathen-powering-potential>

FreeBSD 11 Alpha 1 — New Features Coming To This Open Source OS

FreeBSD foundation has announced the code freeze for the FreeBSD 11 release, marking the start of the Alpha 1 release. This release brings many features, especially for the system administrators and super users.

For those unfamiliar with FreeBSD, it is considered one of the few operating systems left to be true UNIX. It is a direct descendant of the BELL/AT&T labs UNIX. Much of the software available for Linux is also available for FreeBSD, including Gnome and KDE desktop environments and much more user and server software. Despite the amount of software available, it is often thought of as an obscure system with a rather small software library.

This belief is simply untrue, and is less and less true every day due to FreeBSD being able to emulate many Linux system calls. This means much software developed specifically for Linux can be run on FreeBSD as well.

Other operating systems in the same family are Mac OS X, OpenBSD, NetBSD, and Orbis OS, which is the operating system that runs on PlayStation 4 Gaming consoles.

FreeBSD is a longstanding landmark in the open source community, and while it doesn't have as large a user-base as Linux, what it lacks in popularity it surely makes up for in stability, security, and features.

FreeBSD is used in cutting-edge enterprises and is used for the underlying infrastructure of many well-known products like Netflix, WhatsApp, and even as the base for network appliances made by Juniper as well as other companies.

Among the biggest changes to come when FreeBSD 11 arrives in September are updates to:

- The network stack
- Improvements for ARM support
- Several updates to the bhyve hypervisor
- Many security updates
- Storage subsystem updates and an upgraded version of OpenZFS

Are any fossBytes readers FreeBSD users? If so, let us know below in the comments what you think. If not, let us know why not or if you plan on trying it out sometime soon.

<http://fossbytes.com/freebsd-11-release-notes-features/>

Severe Swagger Remote Code Execution flaw compromises NodeJS, Ruby, PHP, Java

This disclosure of an unpatched Remote Code Exec flaw in the Swagger API framework compromises NodeJS, Ruby, PHP, and Java.

Swagger is a representation of RESTful API that allows developers to get interactive documentation, client SDK generation and discoverability.

The Swagger generators are privileged tools for organisations to offer developers easy access to their APIs.

Currently, the Swagger APIs help companies like Apigee, Getty Images, Intuit, LivingSocial, McKesson, Microsoft, Morningstar, and PayPal in building services with RESTful APIs.

Now an unpatched remote code execution vulnerability (CVE-2016-5641) in the Swagger API framework, affecting both client and server components, has been publicly disclosed.

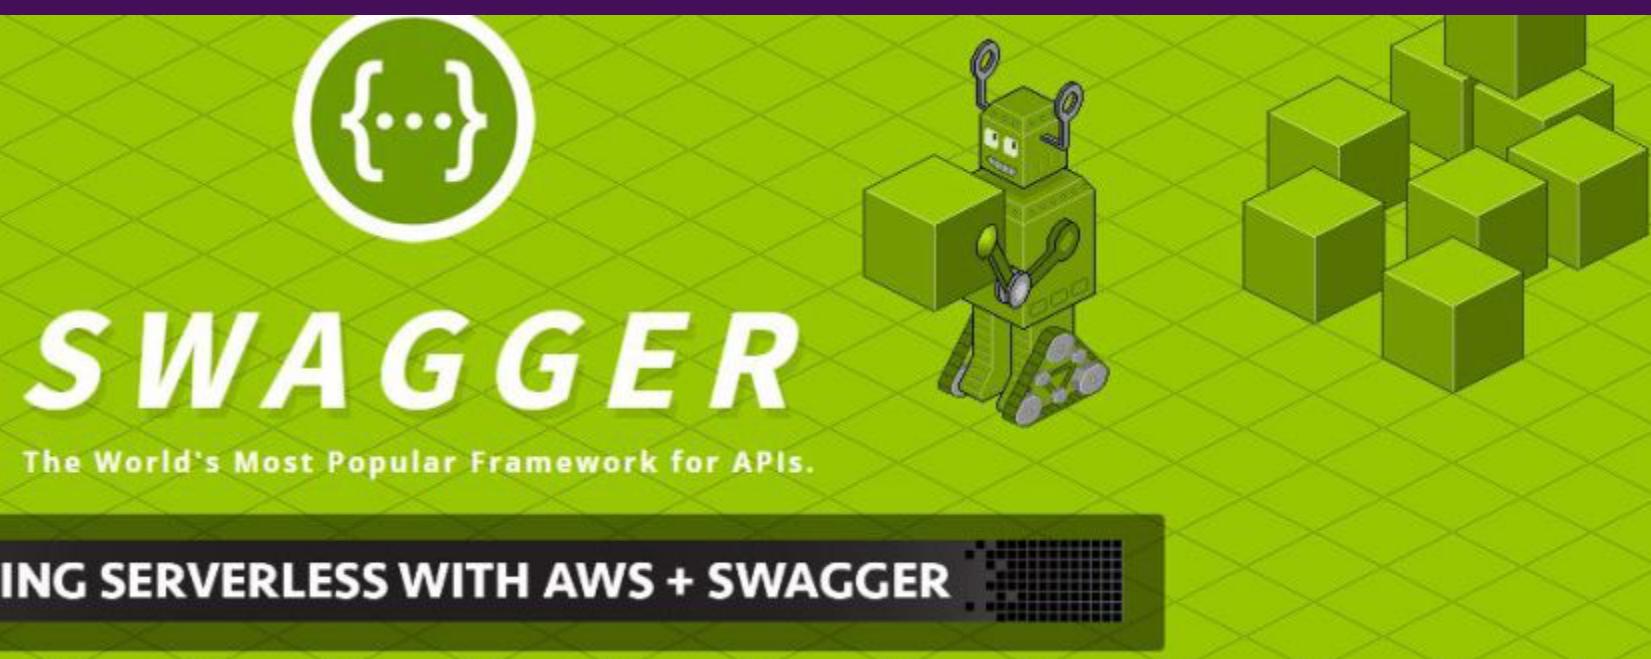
The security vulnerability exists in code generators within the OpenAPI Specification, the REST programming tool.

"The Open API Initiative (OAI) was created by a consortium of forward-looking industry experts who recognize the immense value of standardizing on how REST APIs are described." states the official description.

The remote code execution vulnerability is easy to exploit due to the availability of a Metasploit module released by the security researcher Scott Davis. Davis explained that injectable parameters in Swagger JSON or YAML files allow attackers to remotely execute code across NodeJS, PHP, Ruby, and Java. Davis highlighted that other code generation tools may be vulnerable to parameter injection attacks.

"This disclosure will address a class of vulnerabilities in a Swagger Code Generator in which injectable parameters in a Swagger JSON or YAML file facilitate remote code execution. This vulnerability applies to NodeJS, PHP, Ruby, and Java and probably other languages as well." Davis wrote in a blog post published on the Rapid7 community. "Other code generation tools may also be vulnerable to parameter injection and could be affected by this approach. By leveraging this vulnerability, an attacker can inject arbitrary execution code embedded with a client or server generated automatically to interact with the definition of service. This is considered an abuse of trust in definition of service, and could be an interesting space for further research."

Davis explained that attackers can exploit specially crafted Swagger documents to dynamically create HTTP API clients and servers with embedded arbitrary code execution in the underlying



The image shows the Swagger logo, which consists of a green circle containing three white curly braces (...). Below the logo, the word "SWAGGER" is written in large, bold, white capital letters. Underneath "SWAGGER", the text "The World's Most Popular Framework for APIs." is displayed in a smaller, white, sans-serif font. To the right of the text, there is a 3D-style illustration of a green robot made of cubes, holding a wrench. The background features a light green grid pattern.

GOING SERVERLESS WITH AWS + SWAGGER

operating system. The attack relies on the lack of proper sanitization of the parameters within a Swagger document.

"This is achieved by the fact that some parsers/generators trust insufficiently sanitized parameters within a Swagger document to generate a client code base.

- On the client side, a vulnerability exists in trusting a malicious Swagger document to create any generated code base locally, most often in the form of a dynamically generated API client.
- On the server side, a vulnerability exists in a service that consumes Swagger to dynamically generate and serve API clients, server mocks and testing specs."

The bad news is that the flaw is still unpatched despite it being publicly disclosed; last month the US-CERT issued a specific alert and experts from Rapid 7 already devised a fix.

Rapid7 first attempted to contact the maintainers of the Swagger project in April, exactly one week ago, on June 16, it provided to the US-CERT a patch. The Metasploit module was released on the date of public disclosure, June 23.

Waiting for the release of the patch by the maintainers, users need to carefully inspect Swagger documents for language-specific escape sequences.

<http://securityaffairs.co/wordpress/48679/hacking/swagger-rce-flaw.html>

The current state of open data in the US government

The S.2852 OPEN Government Data Act aims to require true open data access at the federal level. In this article, I will discuss the importance of open data in government, the current state of open data in government, and what we need to do to implement true open data.

When I read an article on the Center for Data Innovation site, Congress Is Stepping Up to Protect Open Data, I was struck by two feelings: elation and surprise. The article explains:

Sponsored by Representatives Derek Kilmer (D-WA) and Blake Farenthold (R-TX) in the House, and Senators Brian Schatz (D-HI) and Ben Sasse (R-NE) in the Senate, the bill would make changes to the U.S. Code to institutionalize open data best practices, such as publishing government data, by default using open and machine readable formats, and with an open license that imposes no restrictions on reuse. "Open by default" has been a mainstay of the open data movement for years, and for good reason: there is simply no way to reliably foresee the potential value of every government data set until the public has access to it.

The article links to a 2013 McKinsey report, Open data: Unlocking innovation and performance with liquid information, that says the national open data initiative of "open by default" represents real economic and transparency value for the public. The report points out that there is no assurance that open data will survive the end of the current administration.

What are the opportunities in legislation for open data within open government?

The how and why of open data was first discussed in the 8 Principles of Open Government Data, published in 2007, and it has increased its importance in government, civil society, and developers' communities. In 2009, President Obama, on his first day in office, issued his first executive order requiring agencies to identify and release "high value" data sets. This had an unintended consequence of some agencies participating and flooding data.gov with what were arguably "high value" data sets and some not participating. In 2013, M1313 Executive Order mandated that all data be open "by default." But did that really happen?

President Obama said of his Executive Order: Making Open and Machine Readable the New Default for Government Information:

"My Administration is committed to creating an unprecedented level of openness in Government. We will work together to ensure the public trust and establish a system of transparency, public participation, and collaboration. Openness will strengthen our democracy and promote efficiency and effectiveness in Government."

One of my open data colleagues, Dr. Dennis D. McDonald, has a concise description of the prob-

lem. He says that unless we address a few questions about open data, the effort will fail:

- How does making open data relate to or support the goals and objectives of the government programs that generate the data?
- How much will open data cost?
- Who will pay?

The executive orders, albeit well intentioned and ground-breaking, did not deliver on the promises made, and executive orders are easy to reverse. Legislation is a stronger approach.

<https://opensource.com/government/16/6/us-open-data-legislation>

Microsoft Releases Open Source .NET Core 1.0 For Linux, Windows, And macOS

Microsoft has released the .NET Core 1.0 for Linux, Windows, and macOS. This announcement adds a new chapter in Microsoft's love affair with Linux and open source. While Red Hat's RHEL Linux distribution currently supports .NET Core 1.0, Microsoft is working to bring the same ability to Debian, Ubuntu, and CentOS.

Back in November 2014, Microsoft promised to make more Core products open source. Now, making good on its promise, Microsoft has finally announced the Windows, Linux, and macOS versions of .NET Core 1.0 and ASP .NET Core 1.0.

The latest open release also includes the release of the .NET Standard Library, making the process of reusing the code across different types of devices like PCs, iOS, and Android much easier.

These announcements were made at Red Hat's tech conference in San Francisco. Notably, Red Hat and Microsoft are often considered fierce rivals in the cloud and enterprise world.

The .NET Core 1.0 is a result of the contribution of more than 18,000 developers and 1,300 companies. Out of these companies, Microsoft's .NET Core 1.0 is officially supported on Red Hat's RHEL Linux distribution and OpenShift.

This partnership with the enterprise-friendly Linux distro will allow enterprises to run different microservices-based apps that are made up of both .NET and Java components.

Microsoft is also working to make .NET Core available for Debian, Ubuntu, and CentOS distribution. The announcement also includes the news that Samsung is joining the .NET Foundation Steering Committee.

"This is the biggest transformation of .NET since its inception and will define .NET for the next decade," Microsoft said.

The same summit also witnessed the release of a new Azure Resource Manager template to make the process of deploying OpenShift on RHEL a simple task.

For more information, you can visit Microsoft's .NET Blog

<http://fossbytes.com/microsoft-releases-open-source-net-core-1-linux-windows-macos/>

iXsystems' TrueNAS Firmware Update Delivers Compelling Performance, Replication, and Graphing Improvements

TrueNAS 9.10 Release Provides up to 40% Improved Performance and More Options to Replicate Data and Graph System Telemetry

SAN JOSE, CA--(Marketwired - July 05, 2016) - Enterprise storage vendor iXsystems today announced the release of Version 9.10 of its TrueNAS enterprise storage array firmware. TrueNAS 9.10 represents the third generation of core software for the award-winning line of enterprise storage arrays and is available to all new and existing TrueNAS users. At the center of TrueNAS 9.10 is an update to the FreeBSD 10.3 operating system which brings significant performance improvements and provides the foundation for the next generation of network fabrics including 100GbE Ethernet.

"My lab benchmarks show TrueNAS 9.10 serves file and block data up to 40% faster than TrueNAS 9.3, especially with higher-end fabrics like 10 and 40GbE. These results should translate to noticeable real-world gains and are just the beginning of what we can do with this solid new platform," said Josh Paetzl, iXsystems Senior Storage Architect.

In addition to performance gains for the SMB, NFS, AFP, iSCSI and Fibre Channel protocols, TrueNAS 9.10 also introduces bidirectional TrueSync replication which allows primary and backup storage hosts to efficiently reverse roles, bringing the other host up to date as needed. This new feature is just one of the many improvements to the TrueNAS High Availability and replication architecture to deliver faster, more resilient live failovers in service of upgrades or outages. TrueSync replication is available between TrueNAS systems on the same LAN or across the globe.

NEWS

For users who actively monitor their storage infrastructure, the TrueNAS 9.10 update expands TrueNAS performance telemetry graphing to include iSCSI and Fibre Channel interfaces, plus the ability to visualize telemetry data on a local Grafana or Graphite server. TrueNAS provides the user with 50 built-in metrics through its web interface and now exposes over 100 metrics to an external server for best-in-class storage system performance analytics.

"TrueNAS 9.10 brings with it a new operating system base that incorporates many new features and performance improvements, all while maintaining its hallmark of stability," says Kris Moore, Director of FreeNAS and TrueNAS Development for iXsystems. "The new FreeBSD 10.3 base will provide the platform for TrueNAS to continue to grow its feature-set in 2016 and beyond."

TrueNAS updates are available through the TrueNAS software updater, a component of the TrueNAS administration interface. TrueNAS users will be notified of the availability of the TrueNAS 9.10 update and should contact iXsystems Technical Support if they have any questions.

To learn more about TrueNAS send an email to sales@ixsystems.com, call 1-855-GREP-4IX, or visit www.ixsystems.com/TrueNAS.

Source: iXsystems Newsletter

Great Specials

On FreeBSD® & PC-BSD® Merchandise

Give us a call & ask about our
SOFTWARE BUNDLES

1.925.240.6652

\$39.95

FreeBSD 9.1 Jewel Case CD Set
or FreeBSD 9.1 DVD

\$29.95

PC-BSD 9.1 DVD

\$49.95

The PC-BSD 9.0 Users Handbook
PC-BSD 9.1 DVD

\$99.95

The FreeBSD CD **or** DVD Bundle

Inside each CD/DVD Bundle, you'll find:

FreeBSD Handbook, 3rd Edition

Users Guide FreeBSD Handbook, 3rd Edition, Admin Guide

FreeBSD 9.1 CD **or** DVD set

FreeBSD Toolkit DVD



FreeBSD 9.1 Jewel Case CD/DVD

CD Set Contains:

- Disc 1 Installation Boot LiveCD (i386)
- Disc 2 Essential Packages Xorg (i386)
- Disc 3 Essential Packages, GNOME2 (i386)
- Disc 4 Essential Packages (i386)

FreeBSD 9.0 CD.....

\$39.95

FreeBSD 9.0 DVD.....

\$39.95

\$39.95

FreeBSD Subscriptions

Save time and \$\$\$ by subscribing to regular updates of FreeBSD

FreeBSD Subscription, start with CD 9.1.....

\$29.95

FreeBSD Subscription, start with DVD 9.1.....

\$29.95

FreeBSD Subscription, start with CD 9.0.....

\$29.95

FreeBSD Subscription, start with DVD 9.0.....

\$29.95

PC-BSD 9.1 DVD (Isotope Edition)

PC-BSD 9.1 DVD.....

\$29.95

PC-BSD Subscription.....

\$19.95

The FreeBSD Handbook

The FreeBSD Handbook, Volume 1 (User Guide).....

\$39.95

The FreeBSD Handbook, Volume 2 (Admin Guide).....

\$39.95

The FreeBSD Handbook Specials

The FreeBSD Handbook, Volume 2 (Both Volumes).....

\$59.95

The FreeBSD Handbook, Both Volumes & FreeBSD 9.1

\$79.95

PC-BSD 9.0 Users Handbook

\$24.95

BSD Magazine

\$11.99

The FreeBSD Toolkit DVD

\$39.95

FreeBSD Mousepad

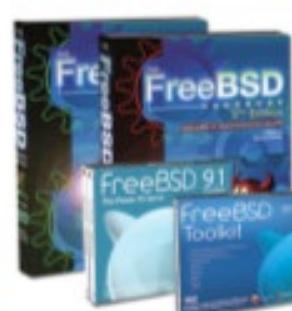
\$10.00

FreeBSD & PCBSD Caps

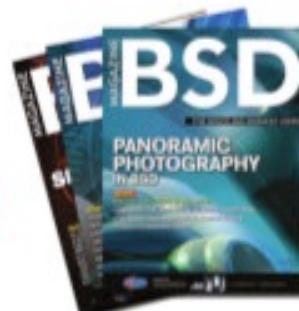
\$20.00

BSD Daemon Horns

\$2.00



Just Plain Fun
Mousepads & Novelty Horns



BSD Magazine
Available Monthly



For even MORE items
visit our website today!

www.FreeBSDMall.com

Bundle Specials!
Save \$\$\$

Implementing In-Memory Cache in the BeaST Architecture

by **Mikhail E. Zakharov**

Recently we have created the reliable, dual-headed FreeBSD based storage concept. Then we improved it with the famous ZFS feature. Now we will try to implement a shared in-memory cache to make our model even more serious.

In the enterprise systems with the real hardware controllers, we should think of connecting the storage cache directly with the high-speed buses that are shared between the controllers. But in our virtual environment, actually, this is still an Oracle VM VirtualBox on my laptop, we have to use LAN for every interconnection. Earlier we decided to stop on the iSCSI protocol for our simple SAN solution.

In general, the virtual environment is not changed much: two machines (ctrl-a and ctrl-b) are the storage controllers and one is the client (clnt-1). But in the previous case with ZFS tests, we had to add four additional virtual drives to emulate SSDs as the storage space for the cache. We do not need them any more because now we intend to use RAM for this purpose. Therefore, we can easily copy our previous virtual machine configurations with the exception of the fxx drives.

The configuration summary is shown in the table below:

Description	ctrl-a	ctrl-b	clnt-1
Inter-controller (private) network. Host-only adapter (vboxnet0)	IP: 192.168.56.10 Mask: 255.255.255.0	IP: 192.168.56.11 Mask: 255.255.255.0	-
Public network. Host-only adapter (vboxnet1)	IP: 192.168.55.10 Mask: 255.255.255.0	IP: 192.168.55.11 Mask: 255.255.255.0	IP: 192.168.55.20 Mask: 255.255.255.0

FreeBSD CORNER

Description	ctrl-a	ctrl-b	clnt-1
Base memory	2048 MB or more	2048 MB or more	Any appropriate value starting with 512 MB will do
Shareable, fixed-sized virtual drives for ZFS data volumes on the SATA controller.	d00, d01, d10, d11 – each drive is 100 MB or more	d00, d01, d10, d11 – each drive is 100 MB or more	-
System virtual drives (Dynamic-sized) on the IDE controller	At least 5 GB to store FreeBSD 10.3-Release default installation	At least 5 GB to store FreeBSD 10.3-Release default installation	At least 5 GB to store FreeBSD 10.3-Release default installation

Note that it's still important to create data-drives (d00, d01, d10, d11) with the fixed-sized and shareable options enabled (see Oracle VM VirtualBox Virtual Media Manager for the details). Otherwise, we will be unable to connect these drives simultaneously to both machines (ctrl-a and ctrl-b) and emulate the drive enclosure that is shared between the storage controllers.

Also, I will continue to use a USB memory-stick to store fixed-sized shareable drives for our testing. But this is being done only in order to slowdown I/O operations and see what is going on in the model.

The latest FreeBSD 10.3 Release can be installed on the non-shareable ada0 drives of the virtual storage machines with the configuration changes in /etc/rc.conf, as usual for our project::

ctrl-a	ctrl-b
hostname="ctrl-a" ifconfig_em0="inet 192.168.56.10 netmask 255.255.255.0" # Inter-controller LAN ifconfig_em1="inet 192.168.55.10 netmask 255.255.255.0" # Public network	hostname="ctrl-b" ifconfig_em0="inet 192.168.56.11 netmask 255.255.255.0" # Inter-controller LAN ifconfig_em1="inet 192.168.55.11 netmask 255.255.255.0" # Public network

FreeBSD CORNER

ctrl-a	ctrl-b
<pre>sshd_enable="YES" # Set dumpdev to "AUTO" to enable crash dumps, "NO" to disable dumpdev="AUTO" # VirtualBox guest additions vboxguest_enable="YES" vboxservice_enable="YES" # iSCSI ctld_enable="YES" # Targets iscsid_enable="YES" # Initiators</pre>	<pre>sshd_enable="YES" # Set dumpdev to "AUTO" to enable crash dumps, "NO" to disable dumpdev="AUTO" # VirtualBox guest additions vboxguest_enable="YES" vboxservice_enable="YES" # iSCSI ctld_enable="YES" # target iscsid_enable="YES" # initiator</pre>

Set iSCSI “disconnection on fail” kernel variable in /etc/sysctl.conf on both systems to be able to failover to the alive controller in case of disaster:

```
kern.iscsi.fail_on_disconnection=1
```

After finishing all these basic FreeBSD installations and preparations, we can start in-memory cache configuration.

The in-memory cache

As we did not write a kernel module with a device driver for our own cache algorithm, we will use ZFS caching capabilities to reach the goal. In the previous article, we have already learned how to configure ZIL and ARC/L2ARC on the shared drives to suit our needs as write- and read-caches. In this part, we will try to allocate these caches in the RAM-drive spaces.

And at this point, we have at least two major problems to solve.

First is that the storage system with its in-memory cache is very sensitive to power loss.

FreeBSD CORNER

Therefore, in real life, we must protect both of the controllers with independent electrical power sources and UPSs. The same requirement is applied to the drive enclosures, as well. Hopefully, now we are playing in the sandbox of the virtual lab without any useful data, so we should not worry about losing it.

Second, we need to invent a synchronization mechanism to enable cache-mirroring feature between the controllers. And in our limited environment, we have to use iSCSI protocol over the virtual LAN as the main transport for this interconnection.

To accomplish this task, we will need to create four memory drives (md) on both controllers to maintain primary and secondary copies of the read- and write- caches and appropriate inter-layer devices – ggate (actually this is ggatel – a “local” version) to simplify ZFS management by using the unified device names. See the table below for the cache structure details:

Memory drive	ZFS inter-layer	Location	Description
md0	ggate0	ctrl-a	ctrl-a write-cache ZFS ZIL primary copy
md1	ggate1	ctrl-a	ctrl-a read-cache ZFS L2ARC primary copy
md2	ggate2 (on gmirror)*	ctrl-a	ctrl-b write-cache ZFS ZIL secondary copy
md3	ggate3 (on gmirror)*	ctrl-a	ctrl-b read-cache ZFS L2ARC secondary copy
md0	ggate0 (on gmirror)*	ctrl-b	ctrl-a write-cache ZFS ZIL secondary copy
md1	ggate1 (on gmirror)*	ctrl-b	ctrl-a read-cache ZFS L2ARC secondary copy
md2	ggate2	ctrl-b	ctrl-b write-cache ZFS ZIL primary copy
md3	ggate3	ctrl-b	ctrl-b read-cache ZFS L2ARC primary copy

* We cant do it right on this step, therefore, see the gmirror configuration on the next page.

FreeBSD CORNER

Now we can implement this structure by the commands:

ctrl-a	ctrl-b
mdconfig -a -t swap -s 128m -u 0	mdconfig -a -t swap -s 128m -u 0
mdconfig -a -t swap -s 128m -u 1	mdconfig -a -t swap -s 128m -u 1
mdconfig -a -t swap -s 128m -u 2	mdconfig -a -t swap -s 128m -u 2
mdconfig -a -t swap -s 128m -u 3	mdconfig -a -t swap -s 128m -u 3
ggate1 create -t 1 -u 0 /dev/md0	ggate1 create -t 1 -u 2 /dev/md2
ggate1 create -t 1 -u 1 /dev/md1	ggate1 create -t 1 -u 3 /dev/md3

We use eight 128 MB memory areas, which gives us 1 GB of cache divided between two controllers. In other words, we have 256 MB of read-cache and 256 MB of write-cache. Actually, these sizes depend on the real memory of the controllers but the basic rule is simple: install more memory and use as much cache as possible.

Now we can establish iSCSI interconnections between cache areas of the opposite controllers. So edit /etc/ctld.conf on both controllers to enable iSCSI-targets. Let's use even LUNs (0 and 2) to access the write-cache and odd LUNs (1 and 3) for the read-cache:

ctrl-a	ctrl-b
portal-group pg0 { discovery-auth-group no-authentication listen 192.168.56.10 }	portal-group pg0 { discovery-auth-group no-authentication listen 192.168.56.11 }

FreeBSD CORNER

ctrl-a	ctrl-b
<pre>target iqn. 2016-01.localsss.private:target0 { auth-group no-authentication portal-group pg0 lun 0 { path /dev/md0 } lun 1 { path /dev/md1 } }</pre>	<pre>target iqn. 2016-01.localsss.private:target0 { auth-group no-authentication portal-group pg0 lun 0 { path /dev/md2 } lun 1 { path /dev/md3 } }</pre>

Tell `ctld` to re-read its target configurations and then initiate iSCSI interconnections:

ctrl-a	ctrl-b
<pre>killall -HUP ctld iscsictl -A -p 192.168.56.11 -t iqn. 2016-01.localsss.private:target0</pre>	<pre>killall -HUP ctld iscsictl -A -p 192.168.56.10 -t iqn. 2016-01.localsss.private:target0</pre>

FreeBSD CORNER

But this was only a half of the job. The other part is to create cache synchronization mechanisms to maintain cache mirroring between the controllers. We also use `gmirror` for this task and `ggate` (`ggatel`) to hide these mirrors under the inter-layer devices (see the remark on the page above):

ctrl-a	ctrl-b
<pre>gmirror load gmirror label ctrl_b_zil /dev/ da0 /dev/md2 gmirror label ctrl_b_arc /dev/ da1 /dev/md3 ggatel create -t 1 -u 2 /dev/ mirror/ctrl_b_zil ggatel create -t 1 -u 3 /dev/ mirror/ctrl_b_arc</pre>	<pre>gmirror load gmirror label ctrl_a_zil /dev/ da0 /dev/md0 gmirror label ctrl_a_arc /dev/ da1 /dev/md1 ggatel create -t 1 -u 0 /dev/ mirror/ctrl_a_zil ggatel create -t 1 -u 1 /dev/ mirror/ctrl_a_arc</pre>

We have just invented the structure of the mirrored devices and their unified inter-layer interface. Now we can try to populate it with the ZFS cache.

ZFS

The Pool configuration is quite simple and well known from the previous article.

We create the `ctrl-a_m0` and the `ctrl-b_m0` pools on the shareable drives (`ada1/ada2` and `ada3/ada4`). Then we allocate volumes on these pools:

ctrl-a	ctrl-b
<pre>zpool create -m none ctrl-a_m0 / dev/ada1 /dev/ada2 zfs create -V 120M ctrl-a_m0/v0</pre>	<pre>zpool create -m none ctrl-b_m0 / dev/ada3 /dev/ada4 zfs create -V 120M ctrl-b_m0/v0</pre>

FreeBSD CORNER

Now we can configure ZFS cache. We will use RAM-drives (hided by the ggate pseudo-devices) for this purpose:

ctrl-a	ctrl-b
# ZIL zpool add -f ctrl-a_m0 log /dev/ggate0 zfs set sync=always ctrl-a_m0 # ARC/L2ARC zpool add ctrl-a_m0 cache /dev/ggate1 zfs set primarycache=none ctrl-a_m0 zfs set secondarycache=all ctrl-a_m0	# ZIL zpool add -f ctrl-b_m0 log /dev/ggate2 zfs set sync=always ctrl-b_m0 # ARC/L2ARC zpool add ctrl-b_m0 cache /dev/ggate3 zfs set primarycache=none ctrl-b_m0 zfs set secondarycache=all ctrl-b_m0

At the final step of ZFS configuration, we will import pools from opposite controllers and set special flags to ignore all errors on them. Yes, this is dangerous but we must do it to be able to fail to the alive controller in case of failure:

ctrl-a	ctrl-b
zpool import -N ctrl-b_m0	zpool import -N ctrl-a_m0
zpool set failmode=continue ctrl-a_m0	zpool set failmode=continue ctrl-a_m0
zpool set failmode=continue ctrl-b_m0	zpool set failmode=continue ctrl-b_m0

FreeBSD CORNER

Note, we set -N option to disable auto-mounting ZFS function.

At this point, ZFS pools of both controllers with their mirrored in-memory caches should work, therefore, we can setup the volume failover arbitrator mechanism.

The Arbitrator

The volume arbitrator, which we use to switch volumes between the controllers, is pretty much the same as in both previous tests.

Therefore, just check and edit /etc/ctld to add appropriate volumes (we use LUN 10 on both controllers for them):

ctrl-a	ctrl-b
<pre>portal-group pg0 { discovery-auth-group no- authentication listen 192.168.56.10 } target iqn. 2016-01.localsss.private:target0 { auth-group no- authentication portal-group pg0 # cache lun 0 { path /dev/md0 } }</pre>	<pre>portal-group pg0 { discovery-auth-group no- authentication listen 192.168.56.11 } target iqn. 2016-01.localsss.private:target0 { auth-group no- authentication portal-group pg0 # cache lun 2 { path /dev/md2 } }</pre>

FreeBSD CORNER

ctrl-a	ctrl-b
<pre>lun 1 { path /dev/md1 } # data volumes lun 10 { path /dev/zvol/ ctrl-a_m0/v0 }</pre>	<pre>lun 3 { path /dev/md3 } # data volumes lun 10 { path /dev/zvol/ ctrl-b_m0/v0 }</pre>

Then, force ctld to re-read its configuration file and reinitialize iSCSI connections to discover the data-volumes. Finally, assemble the arbitrator mechanism by creating active-passive multipath pseudo-device:

ctrl-a	ctrl-b
<pre>killall -HUP ctld # reinitialize iSCSI backend session iscsictl -M -i 1 -p 192.168.56.11 -t iqn. 2016-01.localsss.private:target0 # Create the arbitrator gmultipath create CTRL_B_BACK / dev/da2 /dev/zvol/ctrl-b_m0/v0</pre>	<pre>killall -HUP ctld # reinitialize iSCSI backend session iscsictl -M -i 1 -p 192.168.56.10 -t iqn. 2016-01.localsss.private:target0 # Create the arbitrator gmultipath create CTRL_A_BACK / dev/da2 /dev/zvol/ctrl-a_m0/v0</pre>

FreeBSD CORNER

Note the -M option of the `iscsictl` command that we use to modify the previously established session.

Front-End

In the final part of the storage-side preparations, we should configure the iSCSI front-end to accept the client connections. Therefore, we add “public” related sections to the `/etc/ctl.conf`:

ctrl-a	ctrl-b
<pre>portal-group pg0 { discovery-auth-group no- authentication listen 192.168.56.10 } portal-group pg1 { discovery-auth-group no- authentication listen 192.168.55.10 } target iqn. 2016-01.localsss.private:target0 { auth-group no- authentication portal-group pg0</pre>	<pre>portal-group pg0 { discovery-auth-group no- authentication listen 192.168.56.11 } portal-group pg1 { discovery-auth-group no- authentication listen 192.168.55.11 } target iqn. 2016-01.localsss.private:target0 { auth-group no- authentication portal-group pg0</pre>

FreeBSD CORNER

ctrl-a	ctrl-b
<pre># cache lun 0 { path /dev/md0 } lun 1 { path /dev/md1 } # data volumes lun 10 { path /dev/zvol/ ctrl-a_m0/v0 }</pre>	<pre># cache lun 2 { path /dev/md2 } lun 3 { path /dev/md3 } # data volumes lun 10 { path /dev/zvol/ ctrl-b_m0/v0 }</pre>

FreeBSD CORNER

ctrl-a	ctrl-b
<pre>lun 0 { path /dev/zvol/ ctrl-a_m0/v0 }</pre>	<pre>lun 0 { path /dev/zvol/ ctrl-b_m0/v0 }</pre>
<pre>lun 1 { path /dev/ multipath/CTRL_B_BACK } }</pre>	<pre>lun 1 { path /dev/ multipath/CTRL_A_BACK } }</pre>

Finally, we must tell ctld to update its LUN information once more. Therefore, run on both controllers:

```
# killall -HUP ctld
```

After fresh additions and changes to the concept of the BeaST storage, the structure of our model became quite complex. Therefore, see the full architecture layout below to get the whole picture with a bird's-eye view.

FreeBSD CORNER

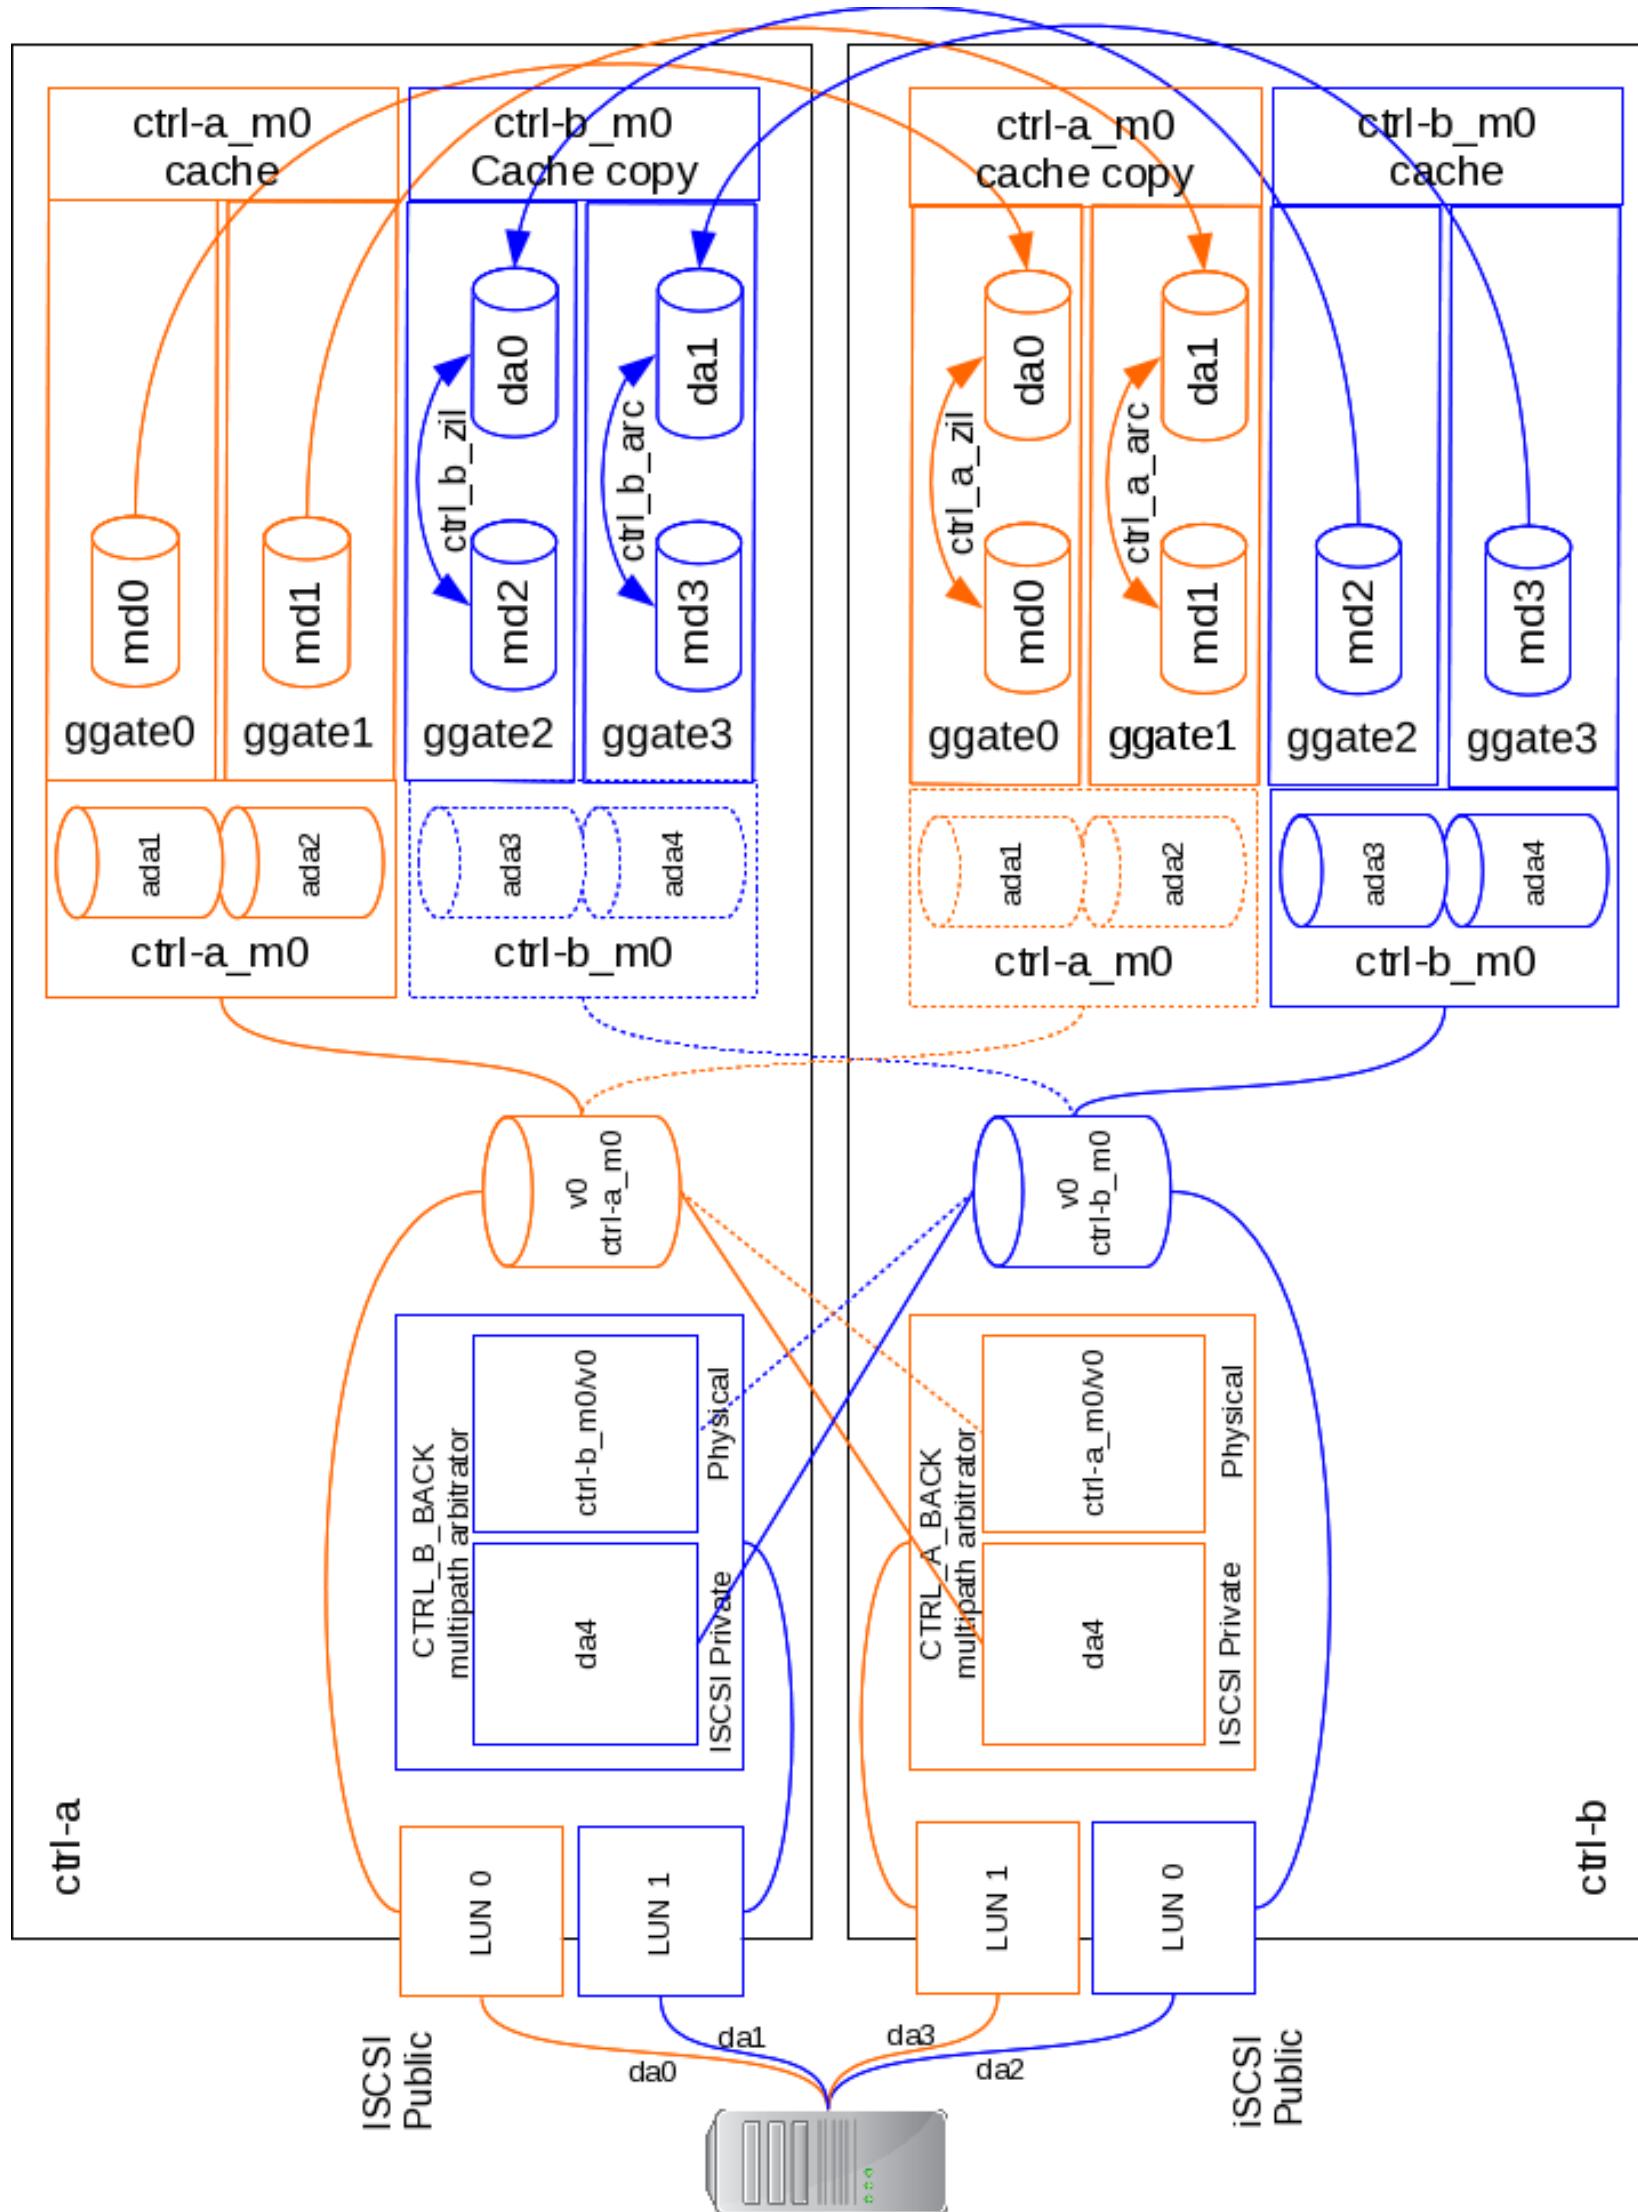


Figure 1. The BeaST storage architecture with ZFS and in-memory cache overview.

FreeBSD CORNER

The client-side

The client configuration has not changed since the first test. So take the previous client virtual machine as-is and check FreeBSD basic configuration in `/etc/rc.conf`:

```
hostname="clnt-1"

ifconfig_em0="inet 192.168.55.20 netmask 255.255.255.0" # Public network

sshd_enable="YES"

# Set dumpdev to "AUTO" to enable crash dumps, "NO" to disable
dumpdev="AUTO"

# VirtualBox guest additions

vboxguest_enable="YES"
vboxservice_enable="YES"

# iSCSI

iscsid_enable="YES"      # Initiators
```

The `/etc/sysctl.conf` file is also the same:

```
kern.iscsi.fail_on_disconnection=1
```

FreeBSD CORNER

Then, run all the very well known commands to bring iSCSI connections online, form multipath pseudo-devices, stripe the volume, create and mount the new file system:

```
root@clnt-1:/home/beast # iscsictl -A -p 192.168.55.10 -t
iqn.2016-01.localsss.public:target0

root@clnt-1:/home/beast # iscsictl -A -p 192.168.55.11 -t
iqn.2016-01.localsss.public:target0

root@clnt-1:/home/beast # gmultipath create CTRL_A /dev/da0 /dev/da3
root@clnt-1:/home/beast # gmultipath create CTRL_B /dev/da2 /dev/da1

root@clnt-1:/home/beast # gstripe create ZDATA /dev/multipath/CTRL_A
/disk/0
root@clnt-1:/home/beast # gstripe create ZDATA /dev/multipath/CTRL_B
/disk/1

root@clnt-1:/home/beast # newfs /dev/stripe/ZDATA
root@clnt-1:/home/beast # mount /dev/stripe/ZDATA /storage
```

Now you can start the common tasks to test the BeaST solution for its functionality and reliability. See our previous works “FreeBSD based dual-controller storage system concept” and “Adding ZFS to the FreeBSD dual-controller storage concept” for the details of these processes.

Conclusion

In this paper, we have offered the mechanism of in-memory cache implementation for the storage system based only on the standard FreeBSD and ZFS features.

This was the most important and interesting part of this paper indeed. So the most common and boring tasks (like FreeBSD installations, CLI outputs along with functional tests, etc.) were cut off from this paper by the Occam's razor principle, as they completely repeat the previous two articles.

We have done much, but there are still more things to do. We must test the BeaST on the real hardware, write installation and clustering scripts to automate typical tasks (like volume creation, migration, recovery, etc) on both controllers. Also, we should check the new and interesting changes in CAM target layer of FreeBSD 10.3.

FreeBSD CORNER

Finally, we must warn you that the BeaST concept is the study of storage systems technology. All the ideas, algorithms and solutions are at concept, development and testing stages. Do not implement the BeaST in production as nobody can guarantee that you will not lose data. Be careful and always do backups!



About the Author:

My name is Mikhail E. Zakharov and I am a proud SAN/storage IBMer. 10 years of experience in large SAN and storage environments: mainly Hitachi, HP and Brocade. Empty – expect-like tool author. FreeBSD enthusiast.

MySQL Server 5.7 Startup Problem on FreeBSD 10.2

by Andrey Ferriyan

I've just installed MySQL server 5.7 for FreeBSD 10.2 using pkg package manager. It seems the problem occurred when I went to start MySQL for the first time. This time I'm using the real VPS server and not VirtualBox.

The error is:

```
/usr/local/etc/rc.d/mysql-server: WARNING:ailed precmd routine for  
mysql
```

Here is what happened to my VPS server.

```
root@cloudfall:~ # service mysql-server start  
/usr/local/etc/rc.d/mysql-server: WARNING: failed precmd routine for mysql  
root@cloudfall:~ # uname -a  
FreeBSD cloudfall 10.2-RELEASE FreeBSD 10.2-RELEASE #0 r296666: Wed Aug 12 16:26  
:37 UTC 2015      root@relenq1.myi.freebsd.org:/usr/obj/usr/src/sys/GENERIC  amd64  
4  
root@cloudfall:~ #
```

FreeBSD CORNER

The default installation from pkg package manager will be in /usr/local/ and the configuration file will be in

```
/usr/local/etc/<your-package>.
```

Some people will try to fix the problem by deleting some line from /usr/local/etc/rc.d/mysql-server. I don't know the result from deleting some line but I had to try using the following method.

1. Fix ownership from my.cnf file in

```
/usr/local/etc/mysql/my.cnf
```

Look at the following figure from my VPS server.

```
root@cloudfall:~ # ls -al /usr/local/etc/mysql/
total 20
drwxr-xr-x  3 root  wheel  512 Jun 21 05:13 .
drwxr-xr-x 19 root  wheel 1024 Jun 21 05:13 ..
drwxr-x---  2 mysql mysql  512 Jun 14 02:01 keyring
-rw-r-----  1 root  wheel 2319 Jun 14 02:01 my.cnf
-rw-r-----  1 root  wheel 2319 Jun 14 02:01 my.cnf.sample
root@cloudfall:~ #
```

Notice that the ownership of my.cnf is user: root and group: wheel. You have to change the ownership to user: mysql and group: mysql. using the following command:

```
# chown mysql:mysql /usr/local/etc/mysql/my.cnf
```

FreeBSD CORNER

2. Next after changing the ownership you can initialize MySQL for the first time before we can use it for production use. The new version of MySQL server uses command `mysqld --initialize` for initialization instead of `mysql_install_db`. If you called only `mysqld --initialize` the command cannot be recognized by the shell because we don't include the file path in the PATH variable. So you have to use an absolute path like this:

```
# /usr/local/libexec/mysqld --initialize
```

3. The following figure will help you understand how the process works.

```
lib/ libdata/ libexec/
root@cloudrain:- # /usr/local/libexec/mysqld --initialize
2016-06-22T06:42:04.770049Z 0 [Warning] Could not increase number of max_open_files to more than 28791 (requested: 32929)
2016-06-22T06:42:04.771526Z 0 [Warning] Changed limits: table_open_cache: 14316
(requested 16384)
 100
 100 200
 100 200
2016-06-22T06:42:05.606586Z 0 [Warning] InnoDB: New log files created, LSN=45790
 100
2016-06-22T06:42:09.142467Z 0 [Warning] InnoDB: Creating foreign key constraint
system tables.
mysqld: Error on delete of './auto.cnf' (Errcode: 2 - No such file or directory)
2016-06-22T06:42:09.234715Z 0 [Warning] World-writable config file './auto.cnf'
has been removed.

2016-06-22T06:42:09.235494Z 0 [Warning] No existing UUID has been found, so we
assume that this is the first time that this server has been started. Generating
a new UUID: 715ee33b-3544-11e6-bab0-0401b5be2201.
2016-06-22T06:42:09.240235Z 0 [Warning] Gtid table is not ready to be used. Table
'mysql.gtid_executed' cannot be opened.
2016-06-22T06:42:10.904890Z 0 [Warning] CA certificate ca.pem is self signed.
2016-06-22T06:42:11.117410Z 1 (Note) A temporary password is generated for root@localhost: E-d|Qpm=aCDDc
root@cloudrain:- #
```

You may notice that from the red square there is a temporary password for managing the MySQL server. But you have to understand that the MySQL server still doesn't start. Now the next step is using the following command for starting the MySQL server:

```
# service mysql-server start
```

FreeBSD CORNER

The following figure below will be the detail information.

```
root@cloudfall:~ # service mysql-server start
Starting mysql.
root@cloudfall:~ # mysql
root@cloudfall:~ # ps -ax | grep mysql
84455 - Ss    0:00.12 /bin/sh /usr/local/bin/mysqld_safe --defaults-extra
85727  = S    0:01.32 /usr/local/libexec/mysqld --defaults-extra-file=/us
85730  1 DL+   0:00.01 grep mysql
root@cloudfall:~ #
```

Do the ps -ax | grep mysql for searching the MySQL server daemon and you can see that MySQL server is up and running.

3. Next write the temporary password from the generator before to /root/.mysql_secret and then do the following command for installing the tables and fix the secure installation.

```
# mysql_secure_installation
```

There will be questions you have to answer. The following figure will help you answer questions from MySQL installation.

FreeBSD CORNER

```
root@bloodfall:~ # mysql_secure_installation
mysql_secure_installation: (ERROR) unknown variable 'prompt=\u001b[1d]\r\_\_'
Securing the MySQL server deployment.

The password file '/root/.mysql_secret' is corrupt! Skipping.
Enter password for user root:

The existing password for the user account root has expired. Please set a new pa-
ssword.

New password:
Re-enter new password:

VALIDATE PASSWORD PLUGIN can be used to test passwords
and improve security. It checks the strength of password
and allows the users to set only those passwords which are
secure enough. Would you like to setup VALIDATE PASSWORD plugin?

Press y|Y for Yes, any other key for No:
Using existing password for root.
Change the password for root ? ((Press y|Y for Yes, any other key for No) :
... skipping.

By default, a MySQL installation has an anonymous user,
```



About the Author:

My name is Andrey Ferriyan. I'm a writer, researcher and practitioner. Python and R enthusiast. Experiences in UNIX-like servers (GNU/Linux, FreeBSD and OpenBSD). Data Scientist wannabe. Area of interests including Information Security, Governance of Enterprise IT, Machine Learning and Data Mining. Now I'm a student at Keio University under LPDP (Indonesia Endowment Fund for Education). I lead startup company in Indonesia called ATSOFT (atsoft.co.id) with my friends. I speak in English and Indonesian.

The article comes from Andreys blog:

<http://andrey.web.id/blog/2016/06/22/mysql-server-5-7-startup-problem-freebsd-10-2/>

Let Us Start at the Very Beginning with Open Source

by E.G. Nadhan

The Red Hat coloring booth at SX Create during the South by Southwest (SXSW) conference and festival helped to tell the open source story of collaboration in a way that was simple enough for children to understand. It was a booth where new ideas manifested themselves in the form of colorful shapes and forms on the walls. Thousands of attendees contributed, bringing the 12 foot tall space to life. Ideas were expanded on by other children, modified or cloned by some to give birth to other related ideas. Continuous, enthusiastic collaboration at its peak in an environment ripe for open communication -- and thus, innovation. Sound familiar?

This is exactly what the Open Source community is all about. The children get it. That is what they do naturally in an uninhibited manner without any artificial boundaries. And thus, the principles of Open Source can be easily injected into the minds of the school-going children. Let us join hands to start at the very beginning with Open Source.

“What I value is the Open Hearts and Open minds”, said Carol Spindel, Lecturer, Dept. of English at the University of Illinois at Urbana Champaign. It was the graduation reception for the Campus Honors Program for the Class of 2016 and Spindel received the Teaching Award for this year. In the brief speech that she gave after receiving this award, Spindel characterized the spirit of the students she works with in her simple but powerful assertion. An Open mind is exactly what Intel VP, Imad Sousou called out in his recent keynote at the Open Stack Summit in Austin. At the same event, I had close encounters with two individuals with fascinating backgrounds and perspectives who bring diversity -- and therefore, an open mindset -- to the world of open source.

Expert Says....

Gina Likins is part of the University Outreach team at Red Hat, which exists to help universities incorporate open source into their curriculum. “In an ideal world, students would graduate from college ready to contribute to an open source community in some significant way” -- reads the abstract for her session at the Open Source Convention in Austin, TX. Likins presented this session along with Professor Heidi Ellis (Western New England University) and Professor Gregory Hislop (Drexel University).

With the Open mindset being a defining characteristic of the student community, the principles of Open Source must be an integral part of the academic curriculum across educational institutions. Here are 6 colleges proactively turning out open source talent. Gina, Heidi, and Gregory also describe the inventory of learning paths they are building to educate students about open source and define a pathway to becoming contributors. This pathway includes a sequence of activities that prepare students for successful humanitarian FOSS participation. The activities cover a variety of aspects of open source participation, including technical, process, and cultural knowledge.

In addition to champions like Gina, Heidi and Gregory, we also need advocates like Charlie Reisinger embedded within the schooling system to drive the adoption of Open Source principles. At Penn Manor School District, Reisinger’s team implemented a unique 1:1 school laptop program using free and open source software exclusively. Reisinger loves to speak on open source learning technologies, and is currently writing a book about their open schoolhouse adventures”. Sidebar: I really enjoyed being part of the #OpenOrg Twitter Chat with Reisinger. Open Minds: Can you take a hint?

Injecting Open Source into academia has a greater likelihood for sparking innovation. Take, for instance, the ninth annual Imagine RIT, the Rochester Institute of Technology’s annual innovation and creativity festival, held this year where about 30,000 people arrive on campus to view student, faculty, and staff demonstrations. Visitors experience everything RIT has to offer through interactive presentations, hands-on demonstrations, exhibitions, and research projects set up throughout campus. One of the tables at this year’s festival was the FOSS@MAGIC table, where students demonstrated their open source software and hardware projects.

There you have it. Forward thinking academic institutions are taking the right steps to fuel the spirit of open source across their student community. And why shouldn’t they? There is a global community of students with the right mindset to embrace, adopt, and innovate the open source way.

The student world is primed and ready for corporate organizations and academic institutions to join forces to inject a healthy dose of Open Source.

Let us start at the very beginning by taking the baby step to introduce Open Source across the schooling system.

Expert Says....



About the Author:

With over 25 years of experience in the IT industry selling, delivering and managing enterprise solutions for global enterprises, E.G.Nadhan is the Chief Technology Strategist at Red Hat (Central Region) working with the executive leadership of enterprises to innovatively drive Cloud Transformation with Open Source technologies and DevOps. Nadhan also provides thought leadership on various concepts including Big Data, Analytics and the Internet of Things (IoT). Nadhan has published 500+ blog posts over four years in multiple Blogs including HP, The Open Group, Enterprise CIO Forum, 1CloudRoad and Intel ITCenter while interacting with analysts from Forbes, Gartner and IDC. Prior to joining Red Hat, Nadhan was an HP Thought Leader and the Global Leader for the Culture of Innovation where he worked with the executive leadership of strategic accounts to realize innovative ideas that address key business challenges for customers. As the Co-Chair for multiple projects within the Cloud Work Group, Nadhan led the publication of the first Cloud Computing Technical Standard and leads the Cloud Governance project within The Open Group. He is a recognized author/speaker at industry conferences and has co-authored multiple books. Follow him on Twitter [@NadhanEG](#). Connect with him on LinkedIn



(<https://www.linkedin.com/in/egnадан>).

Docker Cleanup

by **Nicolas De Loof**

Most Docker newcomers are disappointed when, at some time, after various experiments with Docker, they hit a no space left on device issue. Docker does indeed store all containers, images and volumes in `/var/lib/docker`, which can quickly grow to gigabytes.

One can find tons of blog posts about how to cleanup this folder. Here I'd like to give my own tips, and to explain them in detail, so you don't run random commands found by Googling the Internet.

`/var/lib/docker` is used to store:

- Docker images
- Docker container descriptors
- Docker networks descriptors
- Docker volumes
- Containers' layers (depends on the storage driver you used. Typically AUFS)

Terminated containers

I use Docker a lot to experiment with Unix commands, so I run `docker run --it ubuntu bash`. When you run a container without the `--rm` option, the container still exists on disk after the command completes. So doing this will create tons of containers, and layers for things I modified on the container's filesystem. For this reason, I created an alias `docker-run` to ensure I don't forget this option. But this option can't be used with `-d` (run in background) so I can't use it when I want to run some backend service for testing purpose.

At the end of the day, I have tons of stopped containers that I don't use anymore, and will just consume disk space in /var/lib/docker. So I use this command to run some cleanup:

```
1 | docker rm -v $(docker ps --filter status=exited -q) ?
```

This command lists all exited containers (thanks to status filter), and only dumps their ID (-q). This ID list is used by the remove command to cleanup containers, including volumes they were using.

Unused volumes

Volumes? Yes, when you run a container which has been built with a VOLUME command in it's Dockerfile, Docker implicitly creates a volume on the disk, and may copy data from the container image in this volume. This can result in significant disk consumption. Removing a container with -v option forces Docker to remove such volumes. This doesn't apply to bound mount volumes, only to volumes created by Docker daemon.

If you already removed your containers without using -v, volumes remain orphaned in /var/lib/

```
1 | docker volume rm $(docker volume ls -q -f 'dangling=true') ?
```

docker. You can remove them as well using:

The "dangling" filter selects volumes that aren't referenced by any container.

Note: one can find many scripts to do comparable cleanup by directly making changes to /var/lib/docker. Those scripts were written before docker volume command was introduced in 1.9. Don't use them. Directly hacking your Docker daemon storage isn't a good idea when you have a clean API for this purpose.

Unused images

You can also use a comparable dangling filter with images. This one will detect image layers that

```
1 | docker rmi $(sudo docker images -f "dangling=true" -q) ?
```

are not referenced by a tag, which in many cases is the result of running Docker builds.

What about obsolete / unused images?

You can find some scripts to detect images that aren't used by a container on your system. In a production environment, this can make sense, but for my workstation, this would remove mostly all Docker images, as most containers I run don't keep running all day long.

Docker doesn't track image use by containers, this issue is tracking attempt to change this, so writing a garbage collector would be simpler. So far I'm using <https://github.com/ndeloof/docker-gc> to collect image usage based on Docker events. Not perfect, but does the job.

About the Author:

Nicolas De Loof - Docker Captain, Hacker at CloudBees and community events organizer.

Source of the article - Nicolas' blog:

<http://blog.loof.fr/2016/05/docker-cleanup.html#links>

My Top 9 Favorite Python Deep Learning Libraries

by **Adrian Rosebrock**

So you're interested in deep learning and Convolutional Neural Networks. But where do you start? Which library do you use? There are just so many!

This list is by no means exhaustive, it's simply a list of libraries that I've used in my computer vision career and found particularly useful at one time or another.

Some of these libraries I use more than others — specifically, Keras, mxnet, and sklearn-theano.

Others, I use indirectly, such as Theano and TensorFlow (which libraries like Keras, deepy, and Blocks build upon).

And even others, I use only for very specific tasks (such as nolearn and their Deep Belief Network implementation).

The goal of this blog post is to introduce you to these libraries. I encourage you to read up on each of them individually to determine which one will work best for you in your particular situation.

My Top 9 Favorite Python Deep Learning Libraries

Again, I want to reiterate that this list is by no means exhaustive. Furthermore, since I am a computer vision researcher and actively work in the field, many of these libraries have a strong focus on Convolutional Neural Networks (CNNs).

I've organized this list of deep learning libraries into three parts.

The first part details popular libraries that you may already be familiar with. For each of these libraries, I provide a very general, high-level overview. I then detail some of my likes and dislikes about each library, along with a few appropriate use cases.

Python

The second part dives into my personal favorite deep learning libraries that I use heavily on a regular basis (HINT: Keras, mxnet, and sklearn-theano).

Finally, I provide a “bonus” section for libraries that I have (1) not used in a long time, but still think you may find useful or (2) libraries that I haven’t tried yet, but look interesting.

Let’s go ahead and dive in!

For starters:

1. Caffe

It’s pretty much impossible to mention “deep learning libraries” without bringing up Caffe. In fact, since you’re on this page right now reading up on deep learning libraries, I’m willing to bet that you’ve already heard of Caffe.

So, what is Caffe exactly?

Caffe is a deep learning framework developed by the Berkeley Vision and Learning Center(BVLC). It’s modular. Extremely fast. And it’s used by academics and industry in start-of-the-art applications.

In fact, if you were to go through the most recent deep learning publications (that also provide source code), you’ll more than likely find Caffe models on their associated GitHub repositories.

While Caffe itself isn’t a Python library, it does provide bindings into the Python programming language. We typically use these bindings when actually deploying our network in the wild.

The reason I’ve included Caffe in this list is because it’s used nearly everywhere. You define your model architecture and solver methods in a plaintext, JSON-like file called `.prototxt` configuration files. The Caffe binaries take these `.prototxt` files and train your network. After Caffe is done training, you can take your network and classify new images via Caffe binaries, or better yet, through the Python or MATLAB APIs.

While I love Caffe for its performance (it can process 60 million images per day on a K40 GPU), I don’t like it as much as Keras or mxnet.

The main reason is that constructing an architecture inside the `.prototxt` files can become quite tedious and tiresome. And more to the point, tuning hyperparameters with Caffe can not be (easily) done programmatically! Because of these two reasons, I tend to lean towards libraries that allow me to implement the end-to-end network (including cross-validation and hyperparameter tuning) in a Python-based API.

Python

2. Theano

Let me start by saying that Theano is beautiful. Without Theano, we wouldn't have anywhere near the amount of deep learning libraries (specifically in Python) that we do today. In the same way that without NumPy, we couldn't have SciPy, scikit-learn, and scikit-image, the same can be said about Theano and higher-level abstractions of deep learning.

At the very core, Theano is a Python library used to define, optimize, and evaluate mathematical expressions involving multi-dimensional arrays. Theano accomplishes this via tight integration with NumPy and transparent use of the GPU.

While you can build deep learning networks in Theano, I tend to think of Theano as the building blocks for neural networks, in the same way that NumPy serves as the building blocks for scientific computing. In fact, most of the libraries I mention in this blog post wrap around Theano to make it more convenient and accessible.

Don't get me wrong, I love Theano — I just don't like writing code in Theano.

While not a perfect comparison, building a Convolutional Neural Network in Theano is like writing a custom Support Vector Machine (SVM) in native Python with only a sprinkle of NumPy.

Can you do it?

Sure, absolutely.

Is it worth your time and effort?

Eh, maybe. It depends on how low-level you want to go/your application requires.

Personally, I'd rather use a library like Keras that wraps Theano into a more user-friendly API, in the same way that scikit-learn makes it easier to work with machine learning algorithms.

3. TensorFlow

Similar to Theano, TensorFlow is an open source library for numerical computation using data flow graphs (which is all that a Neural Network really is). Originally developed by the researchers on the Google Brain Team within Google's Machine Intelligence research organization, the library has since been open sourced and made available to the general public.

A primary benefit of TensorFlow (as compared to Theano) is distributed computing, particularly among multiple-GPUs (although this is something Theano is working on).

Python

Other than swapping out the Keras backend to use TensorFlow (rather than Theano), I don't have much experience with the TensorFlow library. Over the next few months, I expect this to change, however.

4. Lasagne

Lasagne is a lightweight library used to construct and train networks in Theano. The key term here is lightweight — it is not meant to be a heavy wrapper around Theano like Keras is. While this leads to your code being more verbose, it does free you from any restraints, while still giving you modular building blocks based on Theano.

Simply put: Lasagne functions as a happy medium between the low-level programming of Theano and the higher-level abstractions of Keras.

My Go-To's:

5. Keras

If I had to pick a favorite deep learning Python library, it would be hard for me to pick between Keras and mxnet — but in the end, I think Keras might win out.

Really, I can't say enough good things about Keras.

Keras is a minimalist, modular neural network library that can use either Theano or TensorFlow as a backend. The primary motivation behind Keras is that you should be able to experiment fast and go from idea to result as quickly as possible.

Architecting networks in Keras feels easy and natural. It includes some of the latest state-of-the-art algorithms for optimizers (Adam, RMSProp), normalization (BatchNorm), and activation layers (PReLU, ELU, LeakyReLU).

Keras also places a heavy focus on Convolutional Neural Networks, something very near to my heart. Whether this was done intentionally or unintentionally, I think this is extremely valuable from a computer vision perspective.

More to the point, you can easily construct both sequence-based networks (where the inputs flow linearly through the network) and graph-based networks (where inputs can “skip” certain layers, only to be concatenated later). This makes implementing more complex network architectures such as GoogLeNet and SqueezeNet much easier.

My only problem with Keras is that it does not support multi-GPU environments for training a network in parallel. This may or may not be a deal breaker for you.

Python

If I want to train a network as fast as possible, then I'll likely use mxnet. But if I'm tuning hyper-parameters, I'm likely to setup four independent experiments with Keras (running on each of my Titan X GPUs) and evaluate the results.

6. mxnet

My second favorite deep learning Python library (again, with a focus on training image classification networks), would undoubtedly be mxnet. While it can take a bit more code to standup a network in mxnet, what it does give you is an incredible number of language bindings (C++, Python, R, JavaScript, etc.)

The mxnet library really shines for distributed computing, allowing you to train your network across multiple CPU/GPU machines, and even in AWS, Azure, and YARN clusters.

Again, it takes a little more code to get an experiment up and running in mxnet (as compared to Keras), but if you're looking to distribute training across multiple GPUs or systems, I would use mxnet.

7. sklearn-theano

There are times where you don't need to train a Convolutional Neural Network end-to-end. Instead, you need to treat the CNN as a feature extractor. This is especially useful in situations where you don't have enough data to train a full CNN from scratch. Instead, just pass your input images through a popular pre-trained architecture such as OverFeat, AlexNet, VGGNet, or GoogLeNet, and extract features from the FC layers (or whichever layer you decide to use).

In short, this is exactly what sklearn-theano allows you to do. You can't train a model from scratch with it — but it's fantastic for treating networks as feature extractors. I tend to use this library as my first stop when evaluating whether a particular problem is suitable for deep learning or not.

8. nolearn

I've used nolearn a few times already on the PyImageSearch blog, mainly when performing some initial GPU experiments on my MacBook Pro and performing deep learning on an Amazon EC2 GPU instance.

While Keras wraps Theano and TensorFlow into a more user-friendly API, nolearn does the same — only for Lasagne. Furthermore, all code in nolearn is compatible with scikit-learn, a huge bonus in my book.

I personally don't use nolearn for Convolutional Neural Networks (CNNs), although you certainly could (I prefer Keras and mxnet for CNNs) — I mainly use nolearn for its implementation of Deep Belief Networks (DBNs).

Python

9. DIGITS

Alright, you got me

DIGITS isn't a true deep learning library (although it is written in Python). DIGITS (Deep Learning GPU Training System) is actually a web application used for training deep learning models in Caffe (although I suppose you could hack the source code to work with a backend other than Caffe, but that sounds like a nightmare).

If you've ever worked with Caffe before, then you know it can be quite tedious to define your `.prototxt` files, generate your image dataset, run your network, and babysit your network training all via your terminal. DIGITS aims to fix this by allowing you to do (most) of these tasks in your browser.

Furthermore, the user interface is excellent, providing you with valuable statistics and graphs as your model trains. I also like that you can easily visualize activation layers of the network for various inputs. Finally, if you have a specific image that you would like to test, you can either upload the image to your DIGITS server or enter the URL of the image and your Caffe model will automatically classify the image and display the result in your browser. Pretty neat!

BONUS:

10. Blocks

I'll be honest, I've never used Blocks before, although I do want to give it a try (hence why I'm including it in this list). Like many of the other libraries in this list, Blocks builds on top of Theano, exposing a much more user friendly API.

11. deepy

If you were to guess which library deepy wraps around, what would your guess be?

That's right, it's Theano.

I remember using deepy awhile ago (during one of its first initial commits), but I haven't touched it in a good 6-8 months. I plan on giving it another try in future blog posts.

12. pylearn2

I feel compelled to include pylearn2 in this list for historical reasons, even though I don't actively use it anymore. PyLearn2 is more than a general machine learning library (similar to scikit-learn in that respect), but also includes implementations of deep learning algorithms.

Python

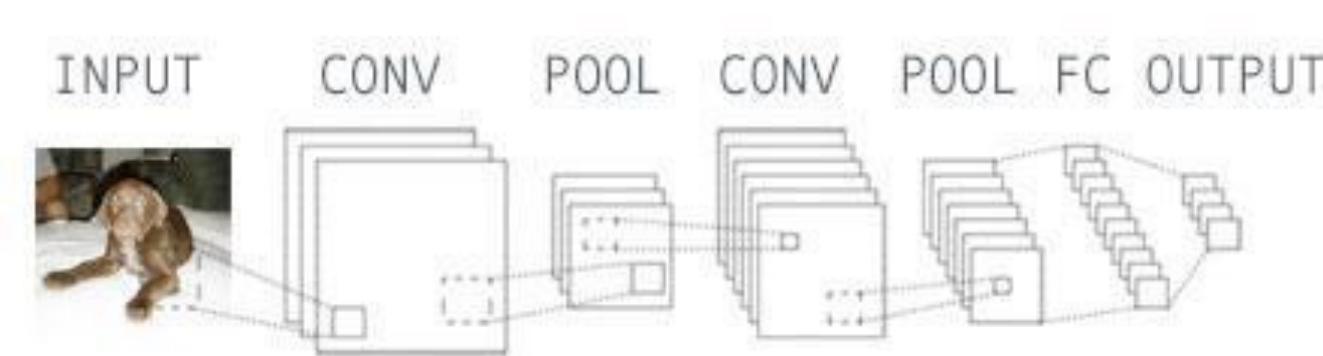
The biggest concern I have with `pylearn2` is that (as of this writing), it does not have an active developer. Because of this, I'm hesitant to recommend `pylearn2` over more maintained and active libraries such as Keras and mxnet.

13. Deeplearning4j

This is supposed to be a Python-based list, but I thought I would include Deeplearning4j in here, mainly out of the immense respect I have for what they are doing — building an open source, distributed deep learning library for the JVM.

If you work in enterprise, you likely have a basement full of servers you use for Hadoop and MapReduce. Maybe you're still using these machines. Maybe you're not.

But what if you could use these same machines to apply deep learning?



It turns out you can — you just need Deeplearning4j.

Take a deep dive into Deep Learning and Convolutional Neural Networks

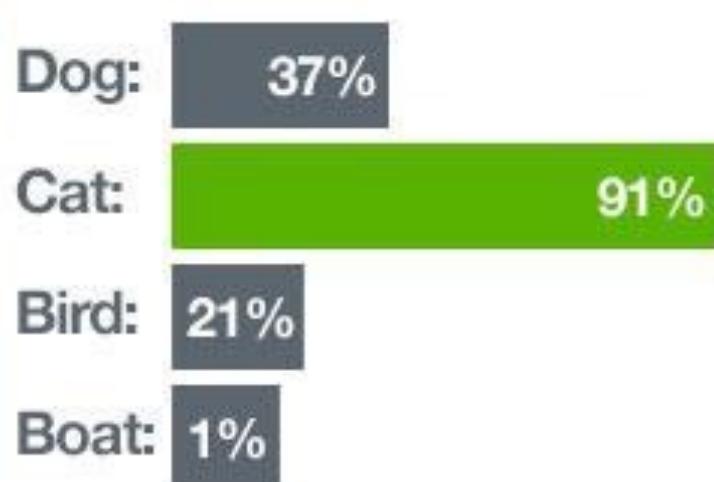
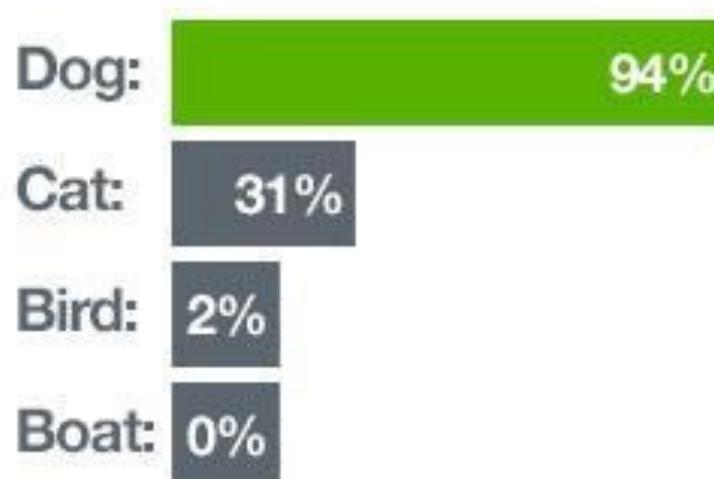


Figure 1: Learn how to utilize Deep Learning and Convolutional Neural Networks to classify the contents of images inside the *Pylmagesearch Gurus* course.

Python

Curious about deep learning?

I'm here to help.

Inside the PyImageSearch Gurus course, I've created 21 lessons covering 256 pages of tutorials on Neural Networks, Deep Belief networks, and Convolutional Neural Networks, allowing you to get up to speed quickly and easily.

To learn more about the PyImageSearch Gurus course (and grab 10 FREE sample lessons).

Summary

In this article, I reviewed some of my favorite libraries for deep learning and Convolutional Neural Networks. This list was by no means exhaustive and was certainly biased towards deep learning libraries that focus on computer vision and Convolutional Neural Networks.

All that said, I do think this is a great list to utilize if you're just getting started in the deep learning field and looking for a library to try out.

In my personal opinion, I find it hard to beat Keras and mxnet. The Keras library sits on top of computational powerhouses such as Theano and TensorFlow, allowing you to construct deep learning architectures in remarkably few lines of Python code.

And while it may take a bit more code to construct and train a network with mxnet, you gain the ability to distribute training across multiple GPUs easily and efficiently. If you're in a multi-GPU system/environment and want to leverage this environment to its full capacity, then definitely give mxnet a try.

Python

About the Author:

My name is Adrian Rosebrock. I've been working in the startup world professionally for the past 8 years, while at the same time earning my PhD in Computer Science with a focus in Computer Vision and Machine Learning at the University of Maryland, Baltimore County.

My Twitter: @PylImageSearch

My Facebook: PylImageSearch

My Google Plus Profile: +Adrian Rosebrock

My LinkedIn: Adrian Rosebrock

You can contact me at adrian (at) pyimagesearch.com.

The article comes from Adrians blog:

<http://www.pyimagesearch.com/2016/06/27/my-top-9-favorite-python-deep-learning-libraries/>

FREENAS MINI STORAGE APPLIANCE

IT SAVES YOUR LIFE.

HOW IMPORTANT IS YOUR DATA?

Years of family photos. Your entire music and movie collection. Office documents you've put hours of work into. Backups for every computer you own. We ask again, *how important is your data?*

NOW IMAGINE LOSING IT ALL

Losing one bit - that's all it takes. One single bit, and your file is gone.

The worst part? **You won't know until you absolutely need that file again.**



Example of one-bit corruption

THE SOLUTION

The FreeNAS Mini has emerged as the clear choice to save your digital life. **No other NAS in its class offers ECC (error correcting code) memory and ZFS bitrot protection to ensure data always reaches disk without corruption and never degrades over time.**

No other NAS combines the inherent data integrity and security of the ZFS filesystem with fast on-disk encryption. No other NAS provides comparable power and flexibility. The FreeNAS Mini is, hands-down, the best home and small office storage appliance you can buy on the market. **When it comes to saving your important data, there simply is no other solution.**

The Mini boasts these state-of-the-art features:

- 8-core 2.4GHz Intel® Atom™ processor
- Up to 16TB of storage capacity
- 16GB of ECC memory (with the option to upgrade to 32GB)
- 2 x 1 Gigabit network controllers
- Remote management port (IPMI)
- Tool-less design; hot swappable drive trays
- FreeNAS installed and configured



<http://www.iXsystems.com/mini>



FREENAS CERTIFIED STORAGE



With over six million downloads, FreeNAS is undisputedly the most popular storage operating system in the world.

Sure, you could build your own FreeNAS system: research every hardware option, order all the parts, wait for everything to ship and arrive, vent at customer service because it *hasn't*, and finally build it yourself while hoping everything fits - only to install the software and discover that the system you spent *days* agonizing over **isn't even compatible**. Or...

MAKE IT EASY ON YOURSELF

As the sponsors and lead developers of the FreeNAS project, iXsystems has combined over 20 years of hardware experience with our FreeNAS expertise to bring you FreeNAS Certified Storage. **We make it easy to enjoy all the benefits of FreeNAS without the headache of building, setting up, configuring, and supporting it yourself.** As one of the leaders in the storage industry, you know that you're getting the best combination of hardware designed for optimal performance with FreeNAS.

Every FreeNAS server we ship is...

- » Custom built and optimized for your use case
- » Installed, configured, tested, and guaranteed to work out of the box
- » Supported by the Silicon Valley team that designed and built it
- » Backed by a 3 years parts and labor limited warranty

As one of the leaders in the storage industry, you know that you're getting the best combination of hardware designed for optimal performance with FreeNAS. **Contact us today for a FREE Risk Elimination Consultation with one of our FreeNAS experts.** Remember, every purchase directly supports the FreeNAS project so we can continue adding features and improvements to the software for years to come. **And really - why would you buy a FreeNAS server from anyone else?**



FreeNAS 1U

- Intel® Xeon® Processor E3-1200v2 Family
- Up to 16TB of storage capacity
- 16GB ECC memory (upgradable to 32GB)
- 2 x 10/100/1000 Gigabit Ethernet controllers
- Redundant power supply

FreeNAS 2U

- 2x Intel® Xeon® Processors E5-2600v2 Family
- Up to 48TB of storage capacity
- 32GB ECC memory (upgradable to 128GB)
- 4 x 1GbE Network interface (Onboard) - (Upgradable to 2 x 10 Gigabit Interface)
- Redundant Power Supply



FreeNAS Getting Started Guide: Part 2, The Initial Configuration Wizard

by Mark VonFange

This article series is intended to serve as an introductory guide to assist FreeNAS users in planning, installation, configuration and administration for their FreeNAS storage systems. This month's article will cover getting your FreeNAS set up using the Initial Configuration Wizard.

The FreeNAS Initial Configuration Wizard

For new users, FreeNAS offers a Wizard system to get your basic functions up and running. Once you log in for the first time, a configuration wizard will pop up to guide you through the initial configuration of your FreeNAS server appliance. If the wizard does not automatically pop up, you can simply click on the top hat icon in the top menu bar. You may quit the Wizard to perform manual configuration any time, but settings are only saved at the end.

You can find full documentation on the Initial Configuration Wizard at https://doc.freenas.org/9.3/freenas_quick.html#initial-configuration-wizard and the FreeNAS team also has put together a video on the Wizard that can be found on their YouTube Channel or on the FreeNAS.org Get Help Video Walkthroughs page at <http://www.freenas.org/about/videos/>.

First, choose your time zone, whether you want to use a language other than English, and whether you will be using a keyboard layout other than the standard US one. Click “Next” to proceed to the next step. These steps may take some time to complete, so there may be a wait before the next screen in the wizard loads.

When getting started, the wizard will first ask you to set up your language, keyboard and timezone settings (Fig. 1).

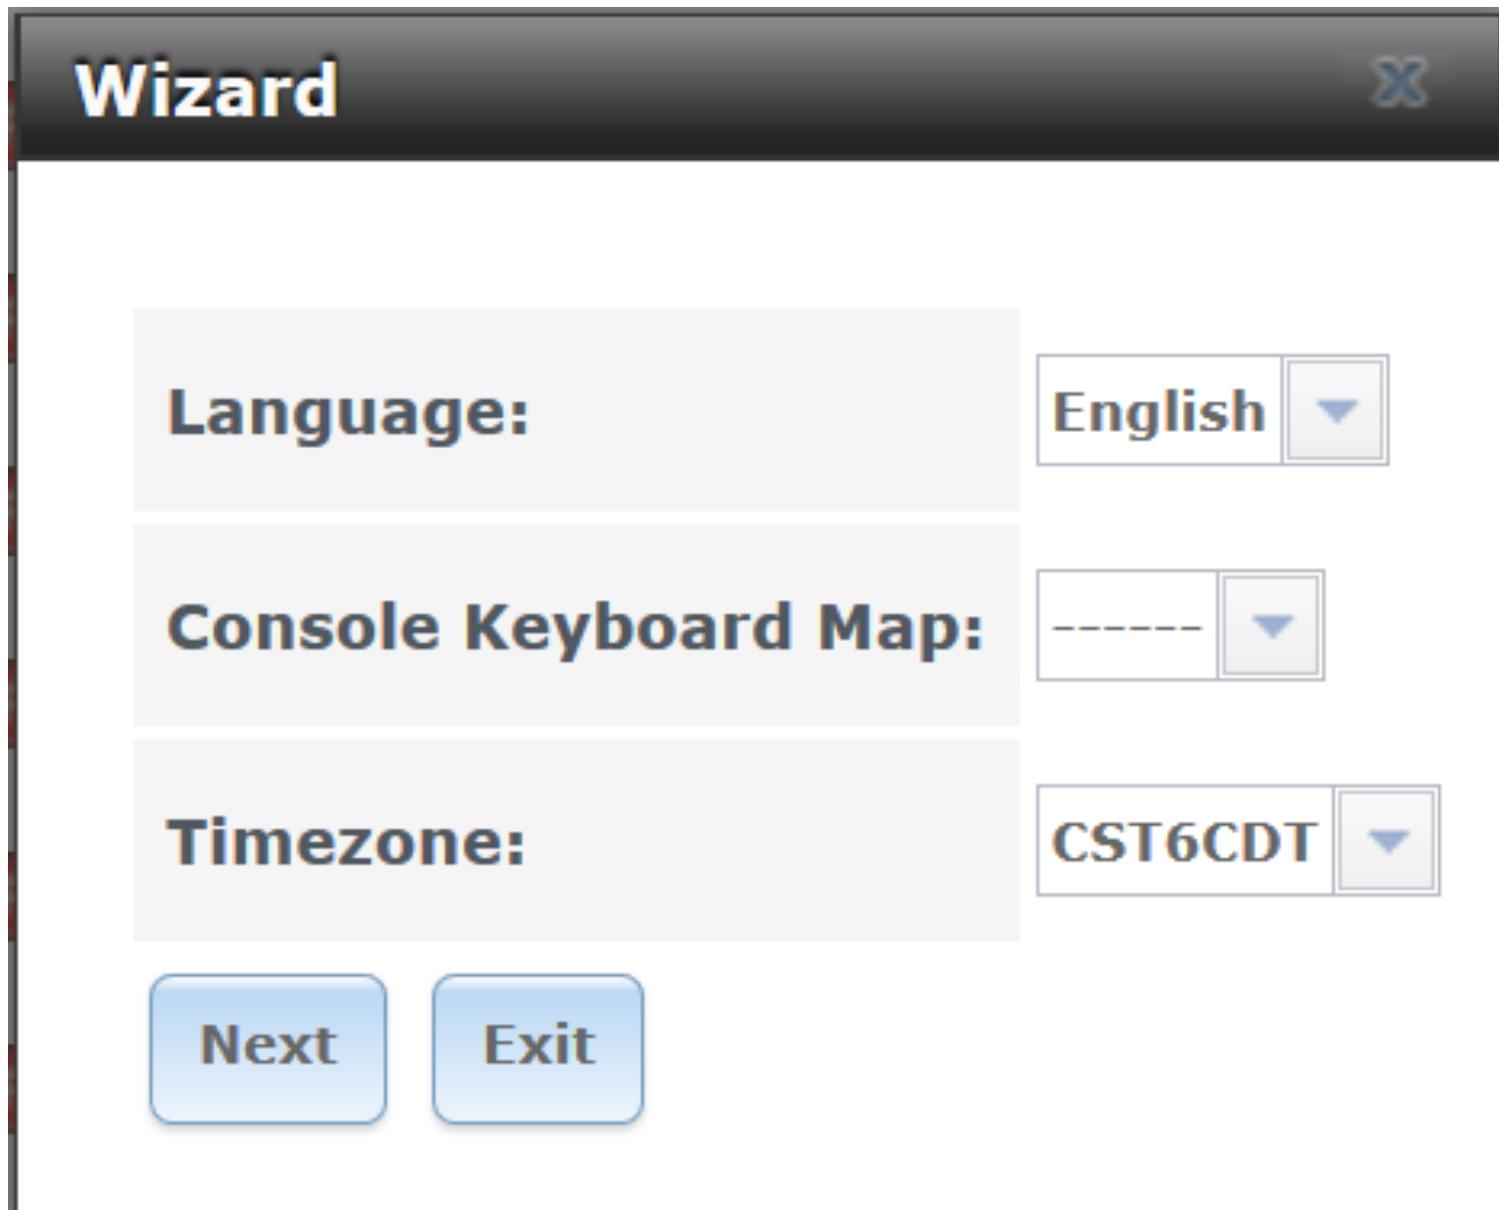


Figure 1. Wizard set up.

- **Volume Configuration**

Next, the Wizard will present some choices for pre-designed volume configurations (Fig. 2). Choose a name for your volume and a standard volume configuration if you wish, then click “Next” to continue. If you have a custom configuration you prefer to use, you may quit the Wizard and perform Volume configuration manually, then click the “Wizard” button to return to the Wizard with the Volume Configuration step left out. You must have a volume configured to continue with the Wizard.

FreeNAS

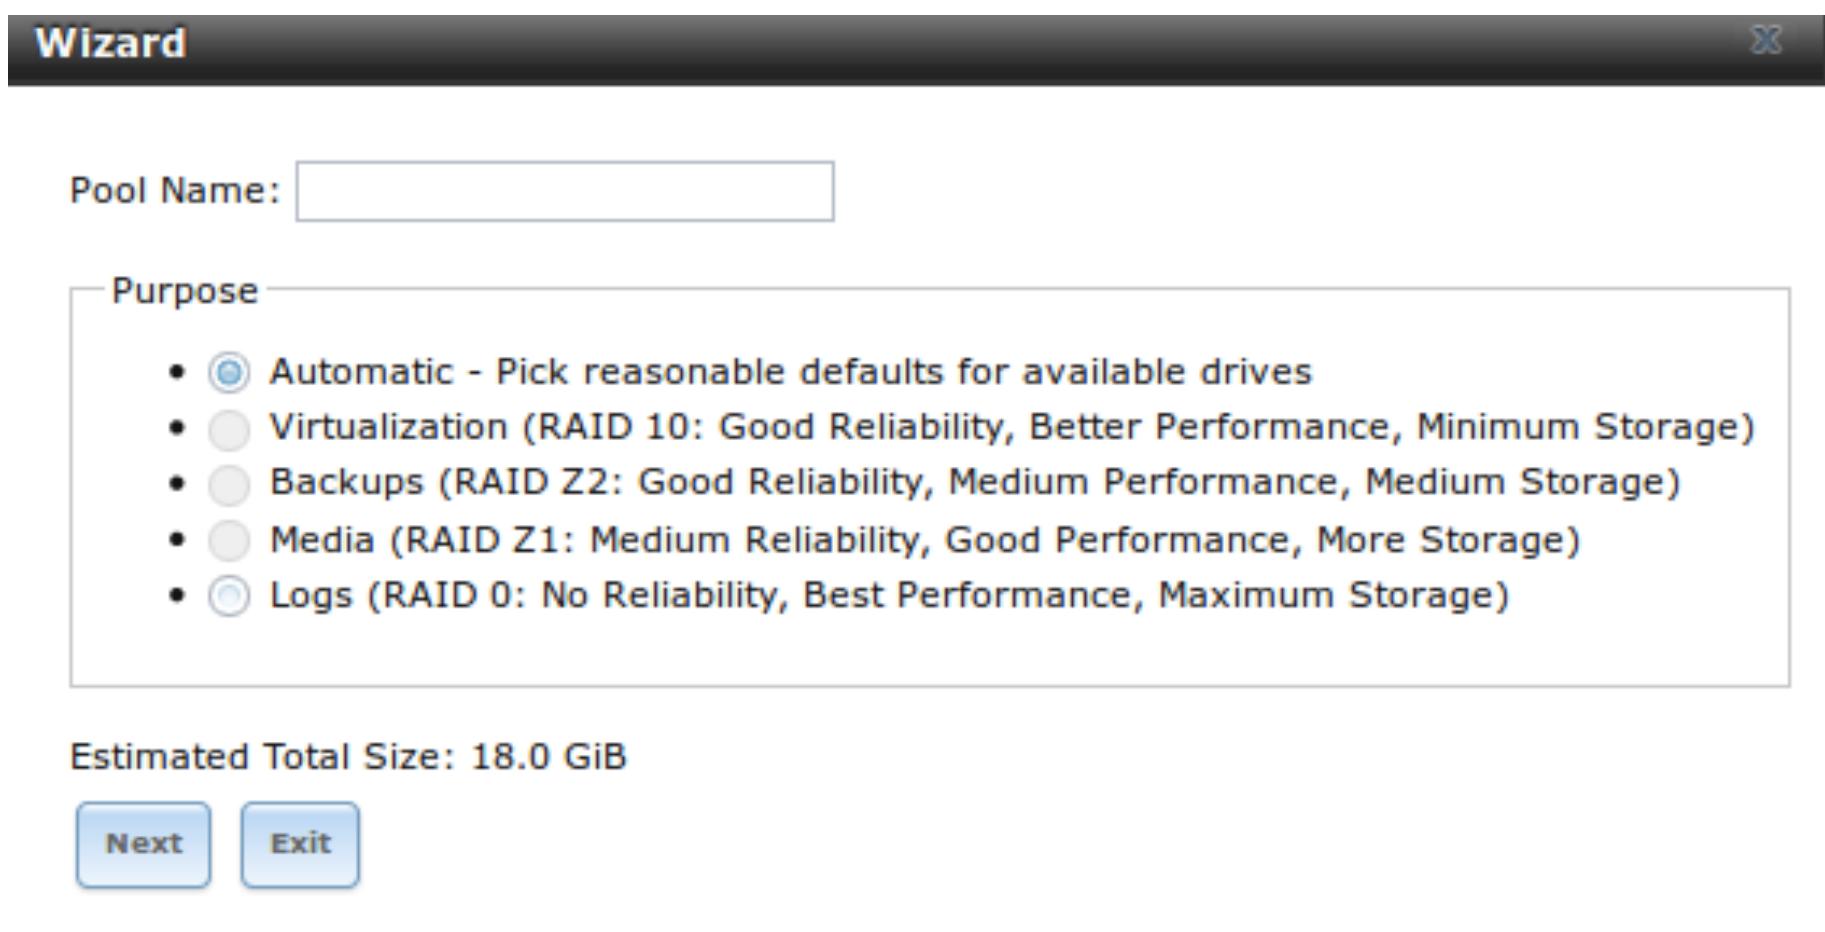


Figure 2: Volume configuration menu.

*Note: If you've already set up your volume, this menu will not appear during the wizard session.

- **Directory Services**

The next menu will allow you to set up basic directory service settings (Fig. 3). If you know your Directory Service account information, select the service you use, enter the applicable information, and then click "Next". If you do not use a directory service, you can click "Next" and proceed to the next step. You can find out more information on this portion in the FreeNAS documentation at http://doc.freenas.org/9.3/freenas_quick.html#initial-configuration-wizard

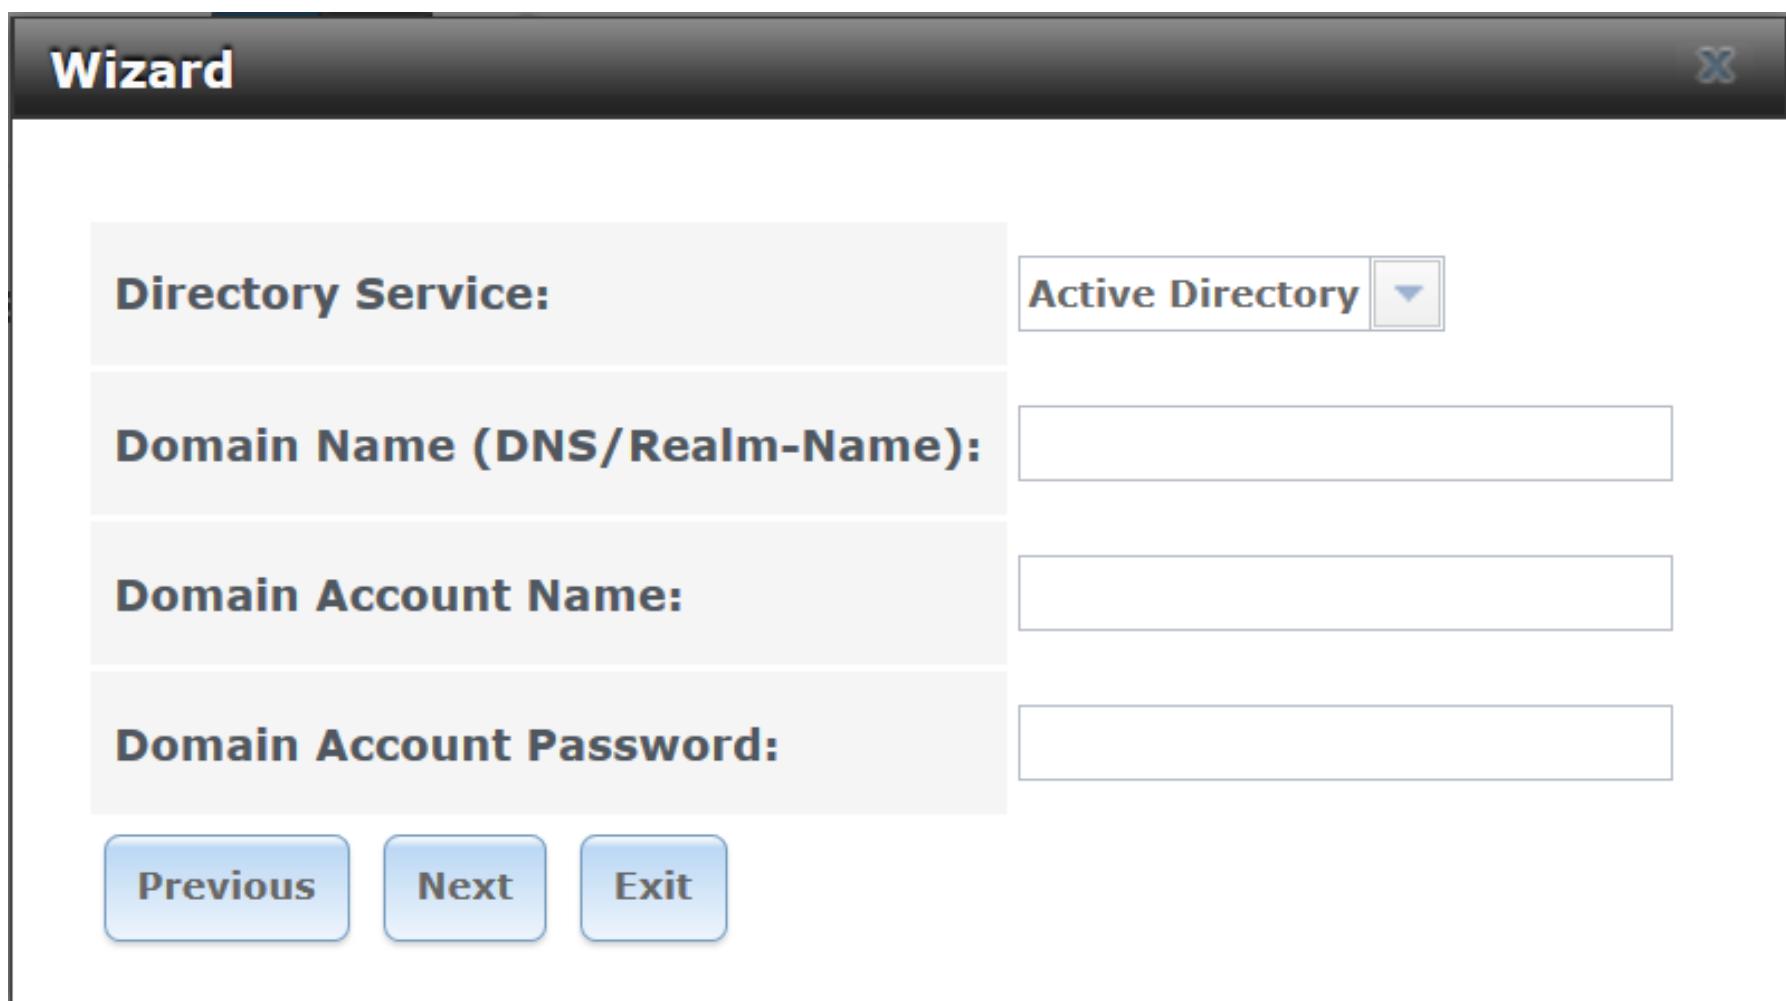


Figure 3: Directory Services Menu

- **Shares and Permissions**

Next, you will need to set up at least one share in order to store files on your FreeNAS server over the network. On this screen (Fig. 4), choose a name for the share. Windows shares are supported by all modern Windows versions, Mac OS X, most popular Linux distributions, and FreeBSD, so we recommend Windows shares if you're not sure what type to pick. If you want the share to be used by anyone on your network without logging in, click the "Allow Guest" button, then click "Add". You may make as many shares as you want, or skip this step if you don't want to add shares at this time.

FreeNAS

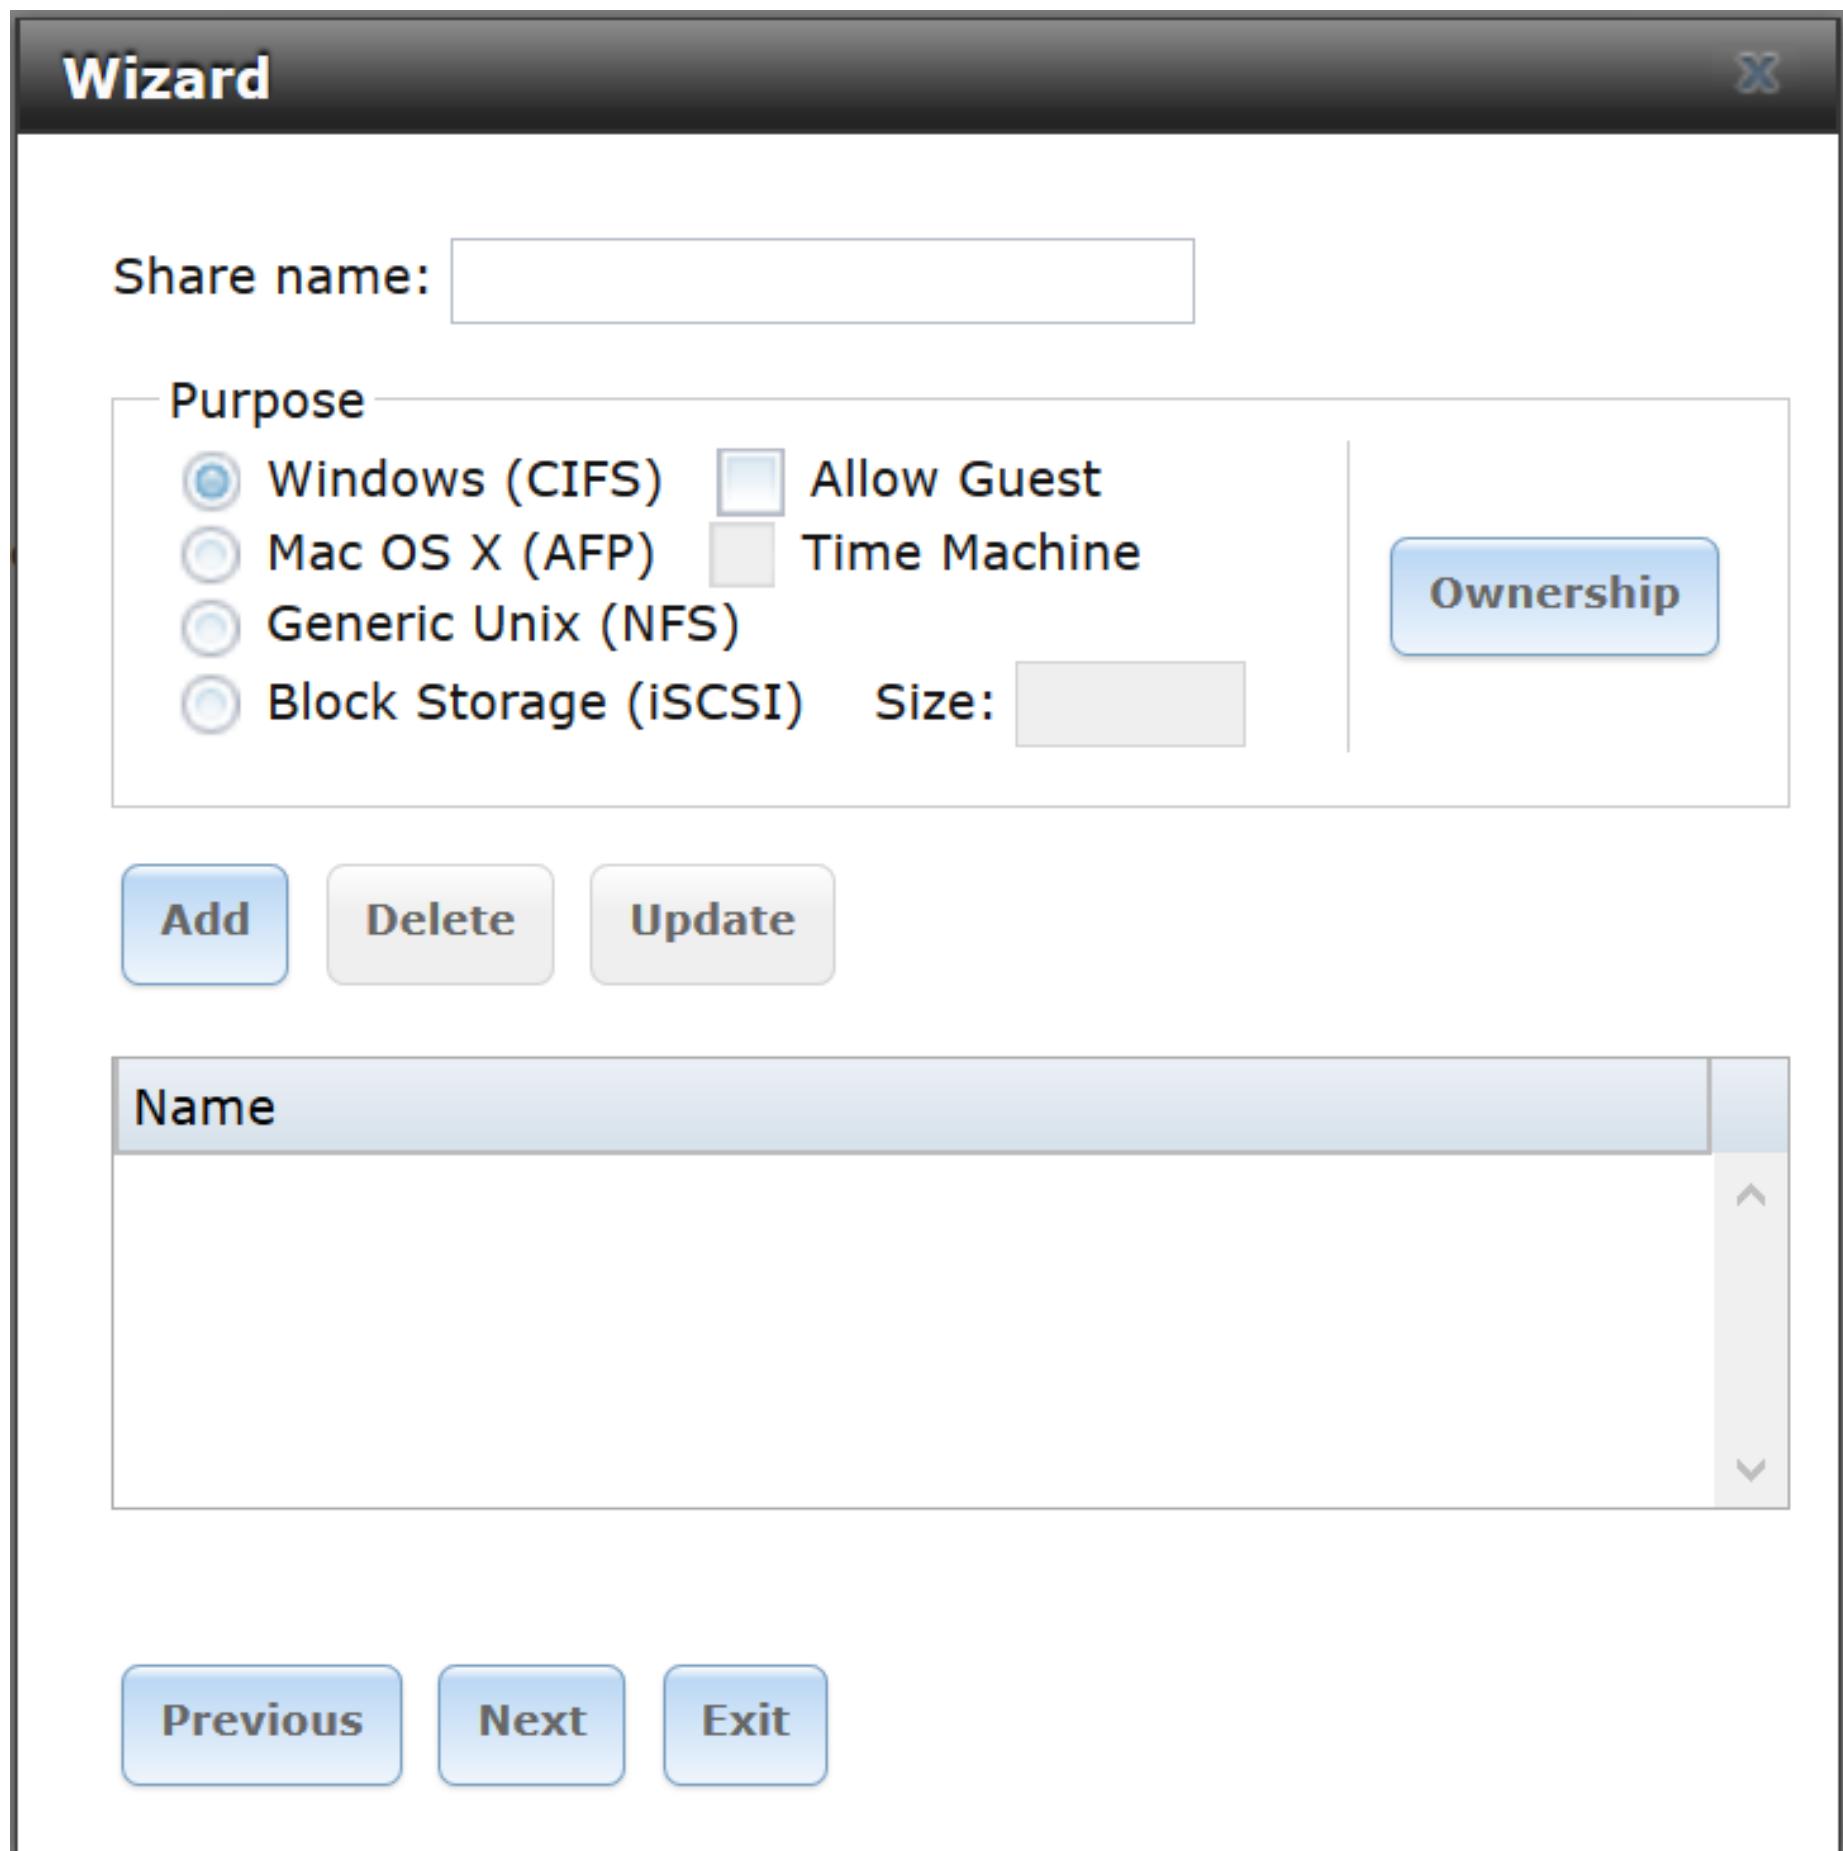


Figure 4: Shares Menu.

FreeNAS

If you instead want to assign a particular owner of the share, click the “Ownership” button. On the Permissions screen (Fig. 5), enter the name of the user you want to own the share and what group you want to share it with. If those aren’t pre-existing system users and groups, click “Create User” and “Create Group” to add them to the system. You will be prompted to create a password to create a new user.

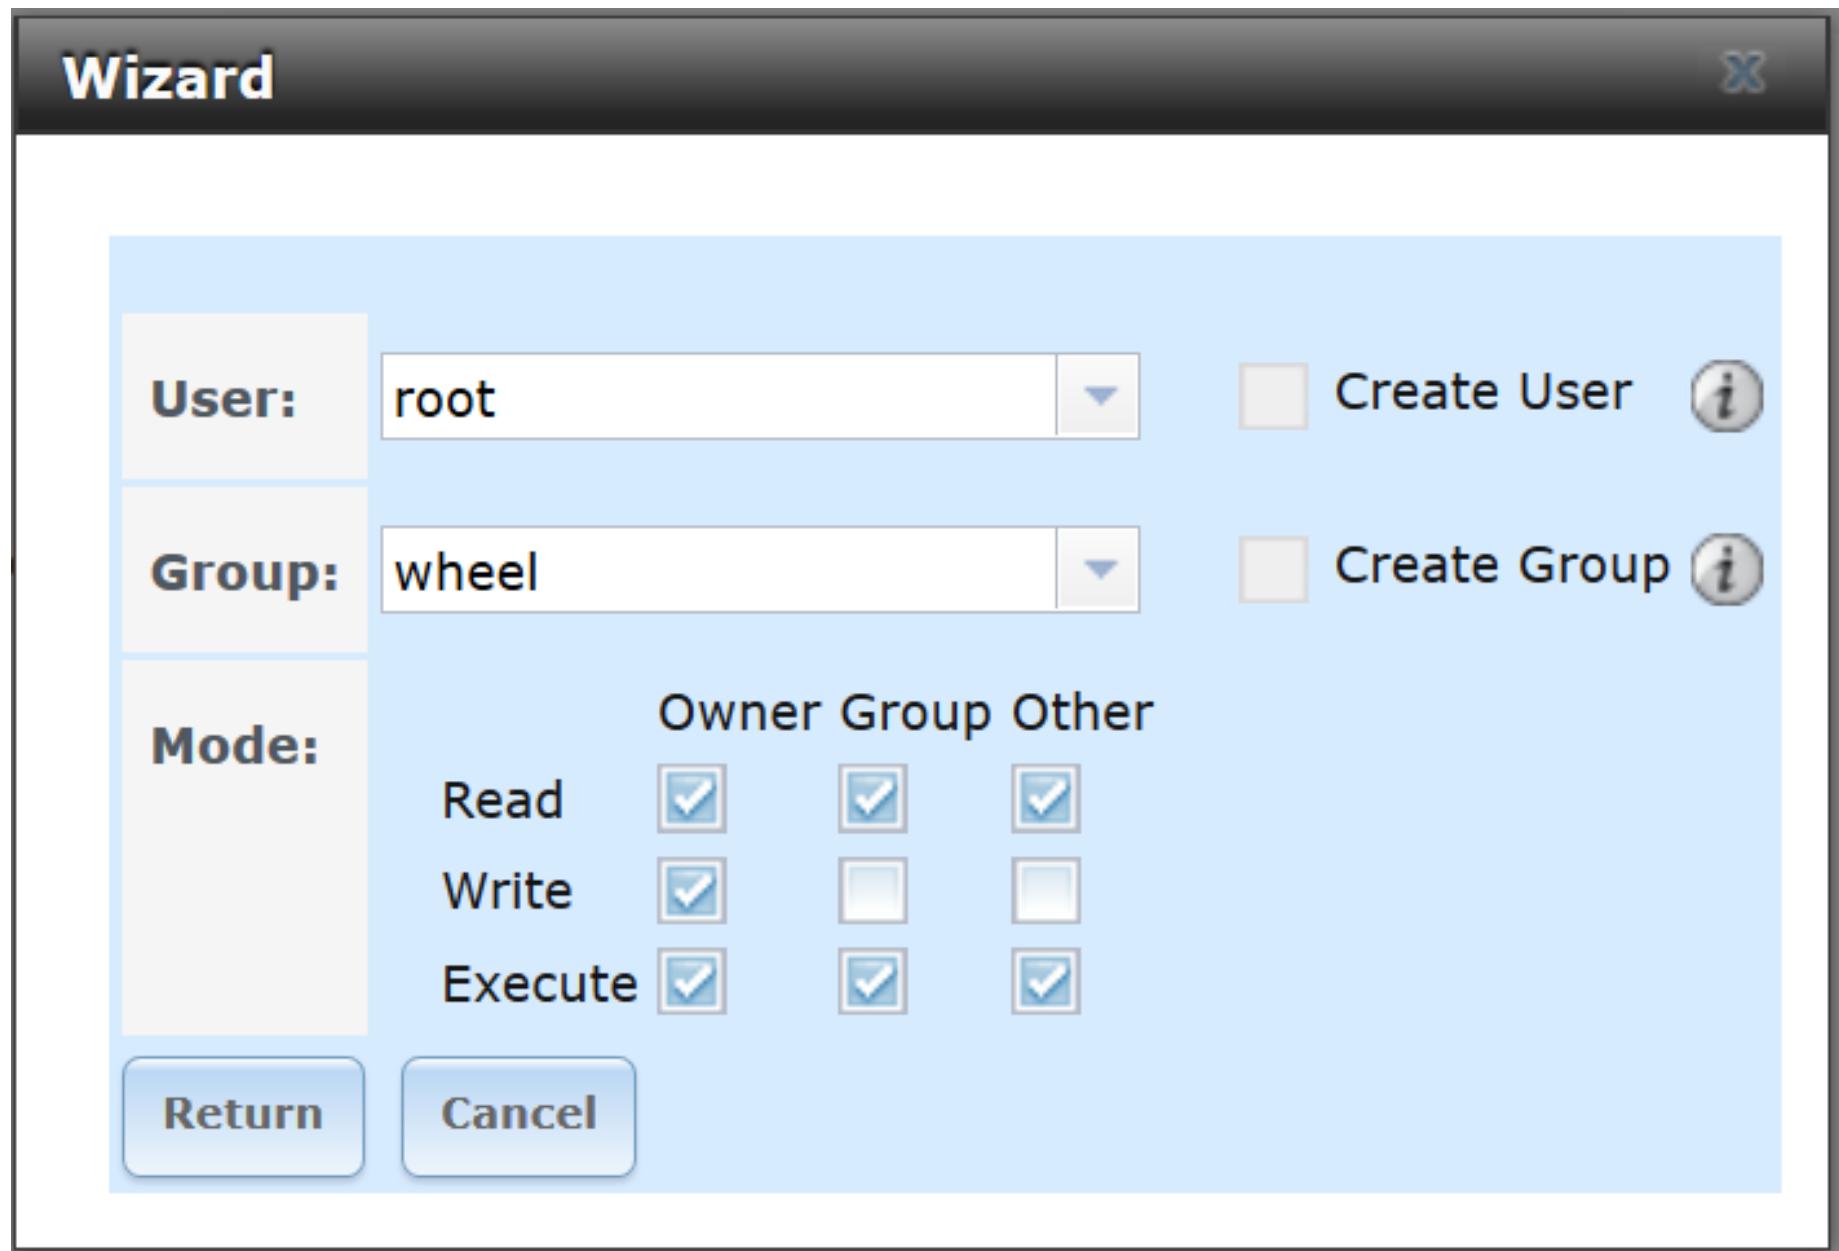


Figure 5: Ownership Menu for User and Group Permissions

You may also choose what permissions are available to the user, group, and other users. After you finish changing the user and group permissions, click “Return” to go back to share configuration. The permissions you selected will apply to the next share you add. You can find out more information on setting up permissions on the FreeNAS YouTube Channel or the FreeNAS.org Get Help Video Walkthroughs page.

*Note: If you’ve already set up a share manually from the User Interface, you will not want to enter anything on this menu due to the potential of data loss.

FreeNAS

- **Miscellaneous Settings**

Next, the Wizard will offer you a number of configuration options (Fig. 6). “Console messages” will enable a live feed of the FreeNAS console log in the footer of the FreeNAS Web Interface. The rest of the settings concern sending emails to alert you of issues with your FreeNAS Mini, such as failed disks or available updates.

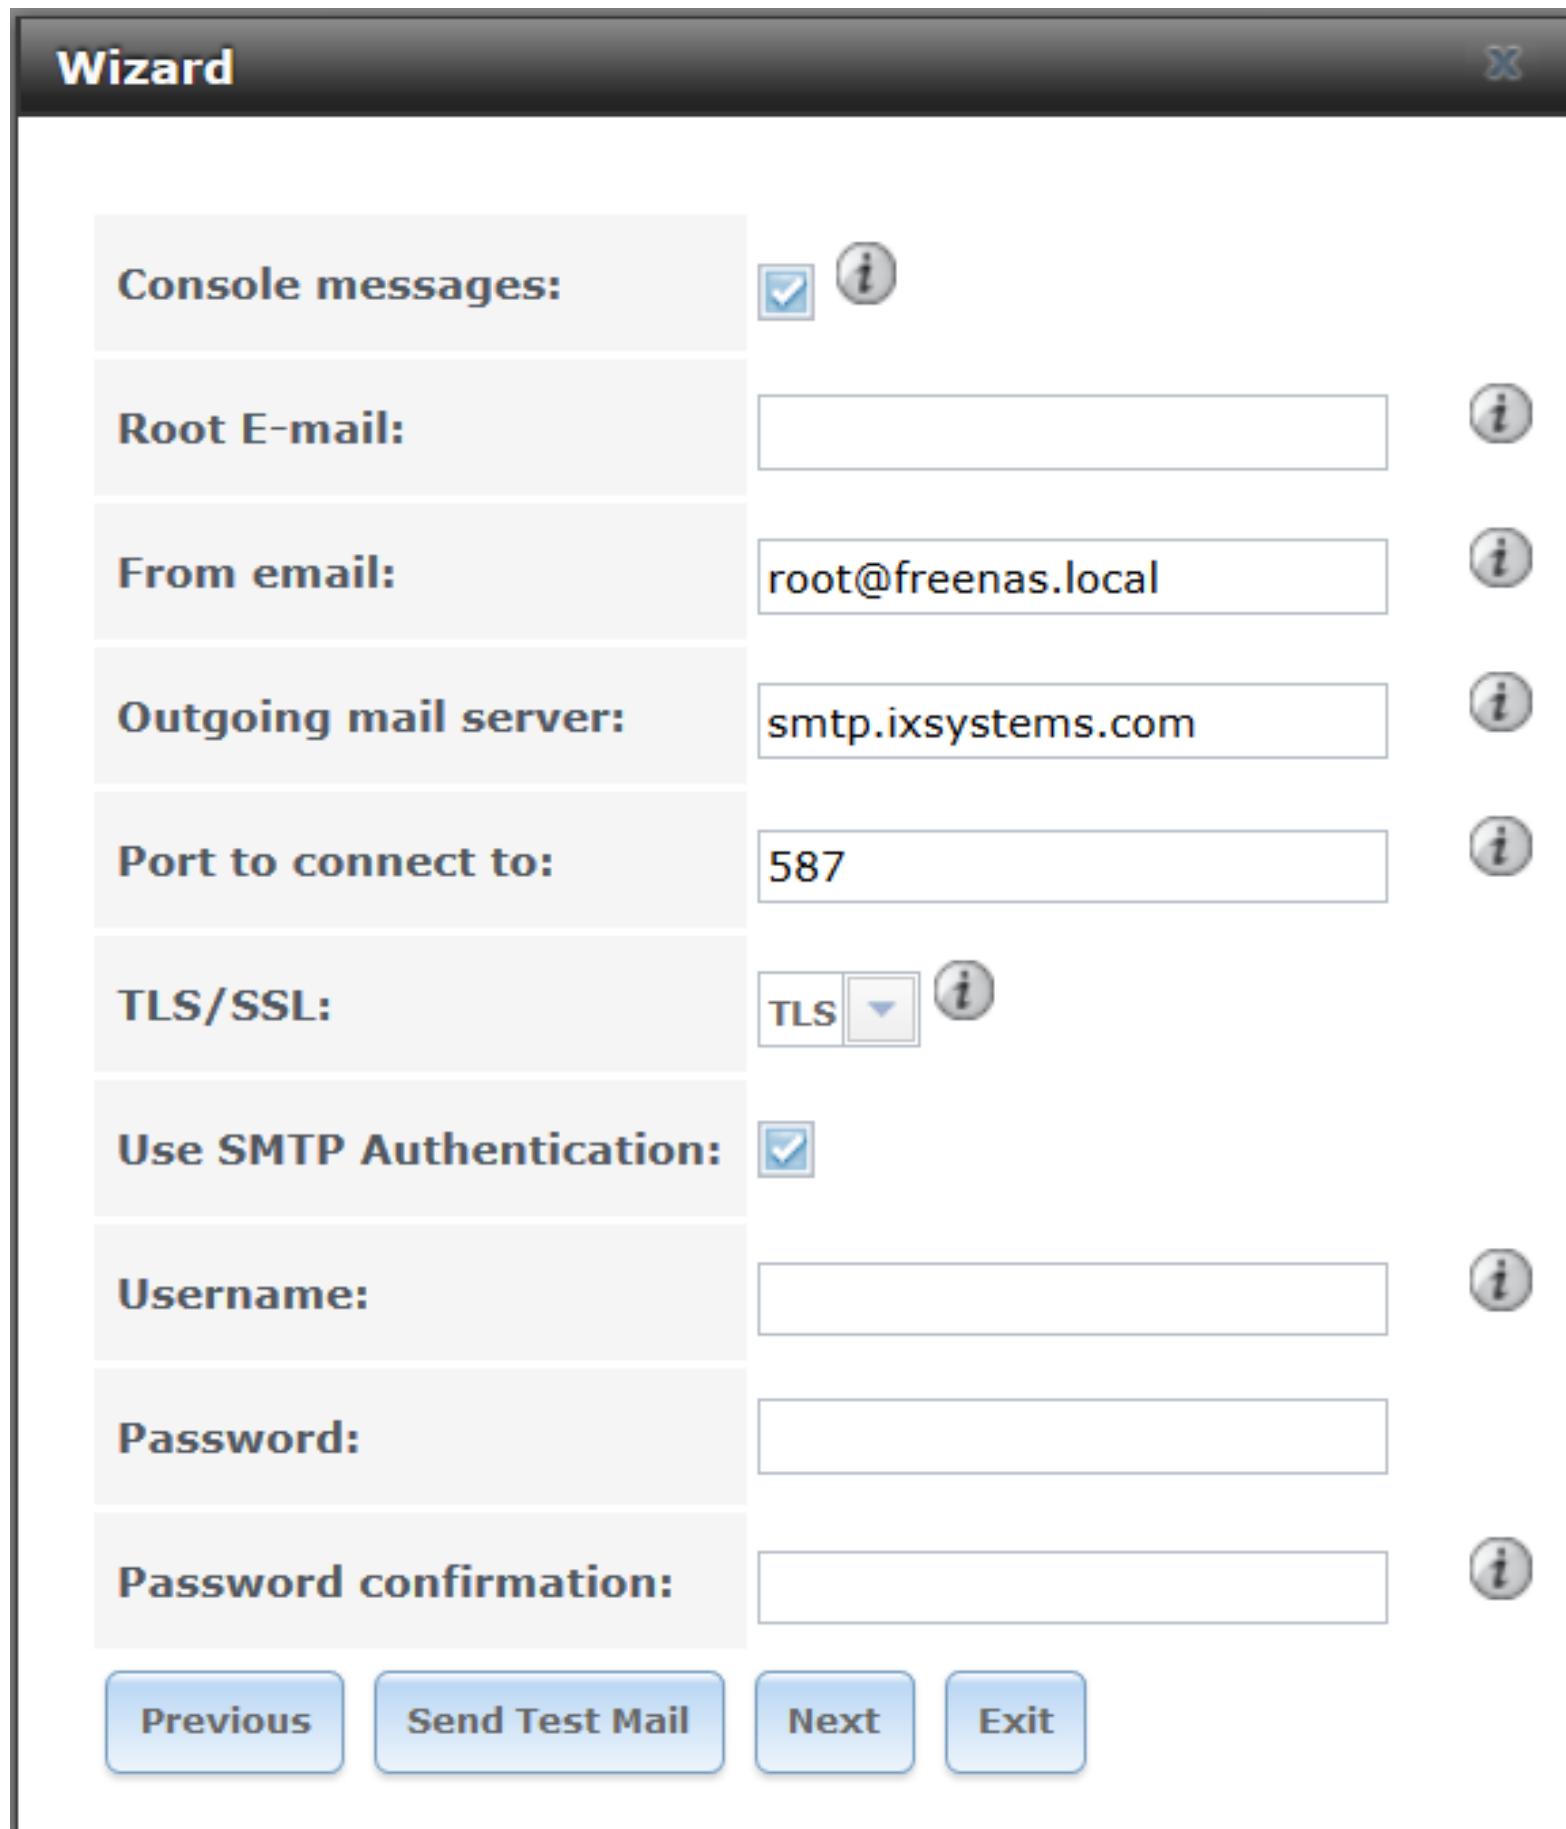


Figure 6: Miscellaneous Settings Menu

Enter the email address into “Root Email” where you would like FreeNAS to send alerts. You must include the hostname or IP address of the SMTP server required by your email account in the “Outgoing Mail Server” field. Port 25 is the default for SMTP, but if you use an alternate port, enter it in “Port to connect to”. If the mail server supports encrypted SMTP connections, select the encryption type in the “TLS/SSL” drop-down menu. If the mail server supports SMTP authentication, check the “Use SMTP Authentication” box and enter the username and password used for mail server authentication.

- **Finish Configuration**

When get to the final Wizard Menu (Fig. 7), if you are sure about all your settings, click “Confirm” to have FreeNAS perform the configuration you’ve selected. If you want to check one more time, click “Return to Wizard”, which will take you back through the steps of the Wizard with all your selected configurations already in place. If you would rather discard your changes and proceed to the FreeNAS UI, you may click “Exit without saving” at this time.

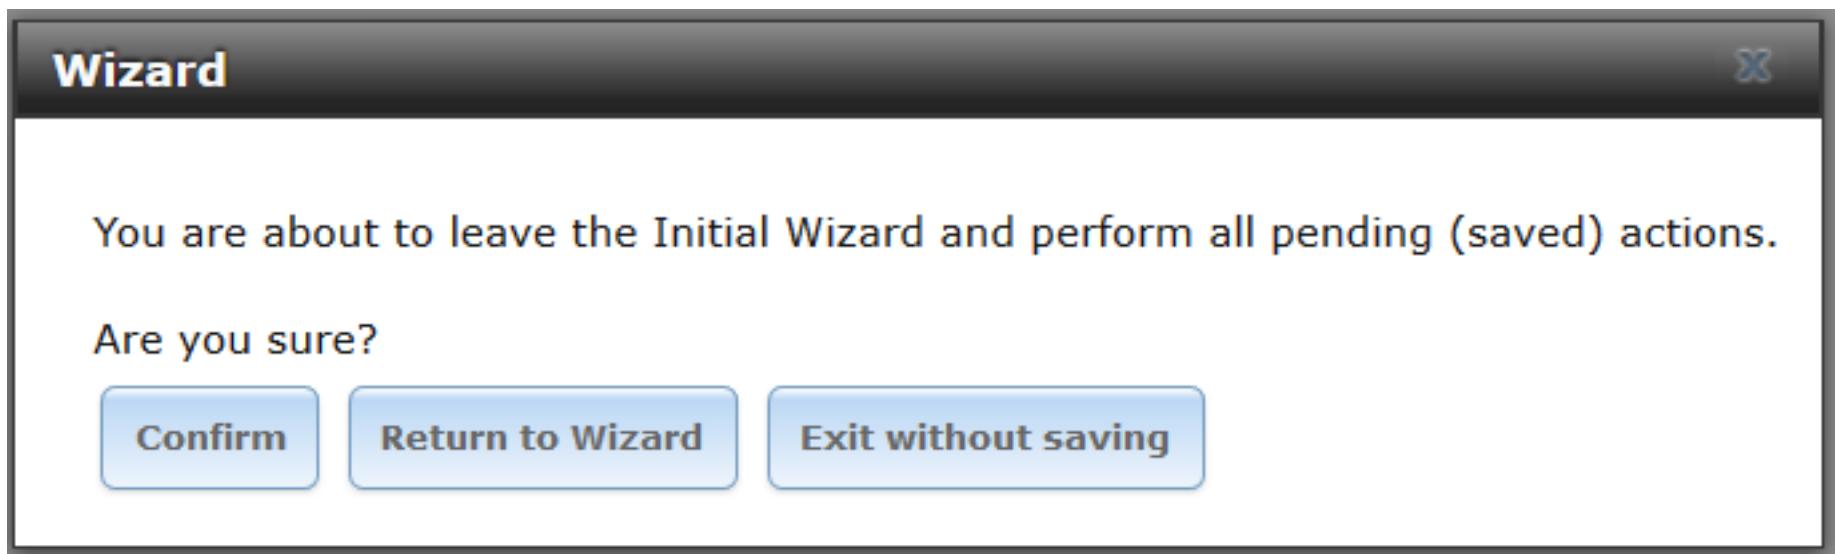


Figure 7. Wizard Menu.

- **Connect to your share**

Now, it is time to connect to your share. The Wizard will already have started the share service(s) you selected, so it is already working. We will use Windows as an example. Open a File Explorer window in a Windows computer on the same network as your FreeNAS server (Fig 8). In the address bar at the top, enter \\ followed by the hostname of your FreeNAS server Appliance (“freenas” by default) or the network IP from your console menu. The current hostname also appears in the title bar of the browser window where you configure your FreeNAS server. Your share contents will appear. You may have to enter the username and password you entered earlier as well. If prompted, enter the username and password you set up in step 5.

FreeNAS

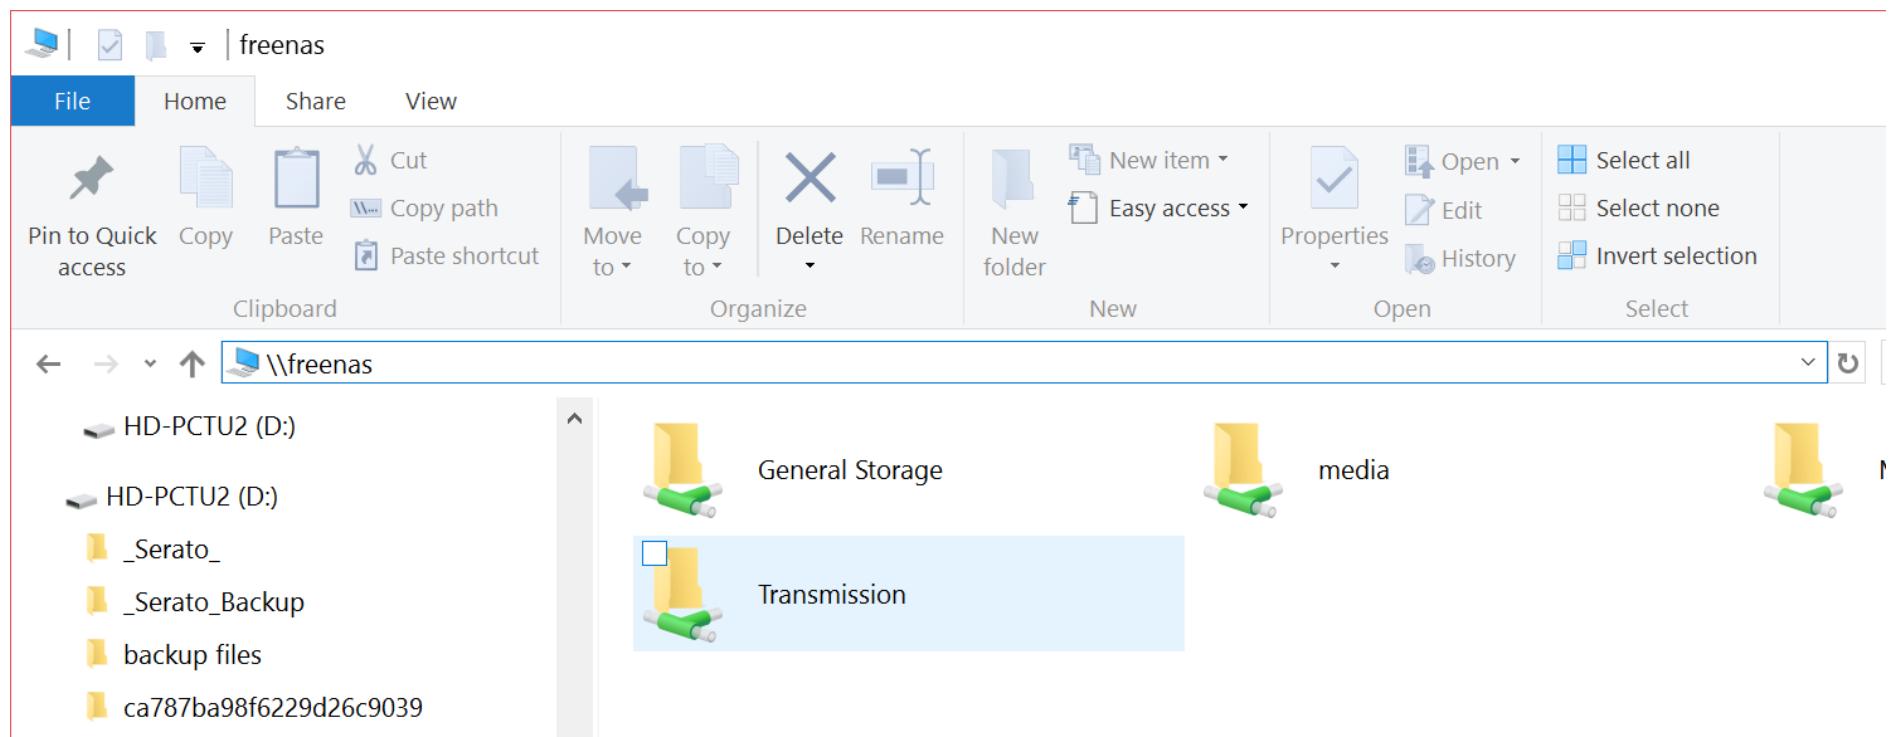


Figure 8: Accessing your FreeNAS server from Windows File Manager

Conclusion

Once you've completed the Initial Configuration Wizard, you should have the core functions of your FreeNAS server up and running for network access. Next month's installment will go over how to accomplish essential tasks from the User Interface itself, like manually setting up users, groups, datasets, shares and services, along with basic administration tasks. Please check out the Additional Resources provided for more FreeNAS related guidance in the meantime.

Additional Resources

Blogs:

- FreeNAS Best Practices: Part 1 | Part 2 | Part 3 | Part 4
(<http://www.freenas.org/blog/a-complete-guide-to-freenas-hardware-design-part-i-purpose-and-best-practices/>)
- FreeNAS: A Worst Practices Guide

Forums: <https://forums.freenas.org/index.php>

Videos:

- Setting up your First NAS with FreeNAS
- FreeNAS 9.x Video Series

Documentation: <https://doc.freenas.org/>

About the Author:

Mark VonFange has worked for iXsystems since 2008 in various roles including first response for professional services inquiries and developing marketing content. He has been an open source advocate for over a decade & enjoys building and repairing computers in his spare time.



USING FREEBSD AS A FILE SERVER WITH ZFS

In this course, we will learn how to use the current ZFS capabilities to help us build a home file server using FREEBSD 10.3.

Course launching date: 04th of July 2016

What will you learn?

- ZFS administration
- ZFS concepts and features

What skills will you gain?

- ZFS administration basics

What do you need?

- FREEBSD 10.3 with root privileges
- At least 10 GB free space

What should you know before they join?

- Basic FREEBSD administration knowledge

WORKSHOP

Module 1: FREEBSD and ZFS

Introduction to ZFS under FREEBSD

- Why ZFS on FREEBSD?
- ZFS features and concepts

Module 2 title: ZFS Administration

Module 2 description: Cover the commands and features to administrate ZFS volumes

- Create, destroy, list pools
- Zpools: single, mirrored, raid
- Understand ZFS properties

Module 3 title: Putting it all to work: Hosting our files using ZFS

Module 3 description: With the previous acquired knowledge, create a plan on how to organize our files and pools to host our files.

- Set ZFS properties based on the content of the files to host
- ZFS tuning
- Create a File Server using our pools

For more info visit our web page:

<https://bsdmag.org/course/using-freebsd-as-a-file-server-with-zfs-2/>

Don't hesitate to ask your questions at

marta.ziemianowicz@bsdmag.org



Do what you love and you will never work a day in your life.

Interview with Aleksandar S. Sokolovski, Co-Founder of Sci Bi

by Marta Ziemianowicz, Marta Strzelec & Marta Sienicka



[BSD Magazine]: Hello Aleksandar, thank you for finding some time for us. How have you been doing? Can you introduce yourself to our readers?

[Aleksandar S. Sokolovski]: Hi I am Alex, Co-founder of Sci Bi (data science solution startup), based in UK (we could expand our business with offices in Canada / USA).

From my alma mater I have a Bachelor of Science in Computer Science and Master of Science in Software Engineering and later on I got my Master's from the University of Sheffield (MSc in TIE - Technology, Innovation and Entrepreneurship). It was a hybrid programme in Computer Science and Business Administration, intended for future manager of tech / service / internet based companies. Later on I got a few certificates, including my pending certificate in Executive Leadership from the University of Cornell (eCornell since it is officially online certification programme). From my work experience, I got my first job one week after my Bachelor's degree in Computer Science as a teaching assistant at a University (I was the first generation of Computer Science students in my country). I worked at the University for about 5 years, then I went to work for a telecommunications and internet provider for about 3-4 years. We officially started our company Sci Bi 2015. From my personal life I am techie, MMO gamer and Sci Fi geek.

[BSD Mag]: Can you tell us something about Sci Bi?

[AS]: The story of Sci Bi is, well, I believe, as unique and interesting as any other startup company founding story. Sci Bi is officially founded by three founders (Me, Dena and Natalie). To tell the story about Sci Bi, I need to start with how the three of us met. I met Dena while I was getting my Master's from the University of Sheffield, we both met with Natalie a few years later.

INTERVIEW

First the story about Dena; as I said, I play a lot of MMO games, most of the games before 2009-2010 were browser based MMO games., While playing one of the war games, my group got into a virtual “war” with another group. We were at virtual war for months and one of the leaders from the group reached out to meet, to discuss peace terms; it was Dena. She messaged everyone from our group leaders and I was the only one who responded to her messages; since she said she was a girl and was sick and tired of waiting for us boys to do anything, she kinda reminded me of my sisters (with her strong female voice not afraid to say and get what she wants, “bossy” is the best word to describe her, it is a positive compliment from my side, since I believe and know females can be great leaders as well as males). After I got talking with Dena, we reached our peace terms for our MMO groups and me and Dena exchanged emails and later on played a few other MMO games where we “rocked” at it. We were even then both data geeks or data natives, and we did some browser scripts to enhance our groups’ game play, we were both very happy when that “primitive” (primitive for now days) script worked. That is when we realized we can do this for real world problems not just games. We exchanged “real names” after a year or so and since she was in a math master program at the same UK university as me, we met in person and worked on a real university project, she has bachelor in mathematics and we did some work together on data analytics. Based on my work with Dena, for my second Master's in Software Engineering, I did my thesis on intrusion detection systems, using data mining and analytics, and after it I never stopped working on data science and big data since (but it was not called data science or big data back in 2010).

Second, the part how Dena and I met Natalie. After our university days, Dena and I stayed in touch as geek friends and we went together to alumni events including one Sci Fi University festival, where we met Natalie. There was some Sci Fi movie, in it there was a female character really not done right as a strong female, but more like a “damsel in distress” character that needs her prince to save her. Natalie started booing at the screen in the middle of the movie screening and I supported that and started booing, as well, but her first reaction was that I was booing at her for being “too female empowerment”, hence Natalie started arguing with me and then Dena got into it as well arguing with her that I was only supporting her argument, not booing against her. Naturally, we got kicked out from the movie screening and we continued our heated argument outside of the cinema, for a few more hours. After it we agreed and disagreed on some points about many sci fi topics including time travel. We agreed to meet again and watch the entire movie and discuss it. We got together online a few times a week to talk about movies and other stuff. Dena and I play Star Wars The Old Republic MMO game, we even got Natalie to play and it she liked it.

Games, multiplayer, MMO and the “online” world are a great part of our company culture. I met one of Sci Bi’s lead software engineer at an E3 game convention; he challenged me to a game of Dragon Age Inq. He was victorious, and we are great friends and a good team now. Dena met Lyhn (she is in our Sales team) at a game tournament for Heroes of the Storm.

INTERVIEW

In the past month or so at the office when we stay to work over the weekend, our favorite game is Star Wars Battlefront and Heroes of the Storm, also a few racing games like Need for Speed, and one on one fighting, like Mortal Kombat.

The last part is the story of how Sci Bi started. When Star Trek Into Darkness got released in 2013, we got together to watch it and discuss it, and we all agreed that the female science character was classical “damsel in distress” (but JJ did great work on Star Wars VII with Ray). We realized that for any good “hero” in a movie you need a great “villain”. The same can be said for startups; for a good startup you need a great challenge, and data science was that challenge for us. After one year, we found a VC support and started our company officially in early 2015 (January/February). The three of us are a great team since we complement each other. Natalie is marketing / sales lead, Dena is COO material from start, and I am the tech geek who likes to “play” with new tech. There was a TV show “House” where the doctor used brainstorming as the main source and method for problem solving, plus he never liked working with someone who thinks like him and agrees with him on every topic and point. I am kinda like that, I need someone to challenge my ideas not always 100% agree with them.

The three of us used brainstorming daily even if we are not in the same country / continent. A few months back, while I was traveling for business conference to Vienna (btw Vienna in one of my favorite cities in Europe, along with Warsaw and Frankfurt), we used slack to brainstorm about our new AI service with Dena and Natalie. Dena was, I believe, at the Cape Town University (that is in South Africa) for a science conference and Natalie was in New York for a marketing conference. Natalie was officially the CEO for the past two years, I was recently appointed CEO a few months back, but it was because Natalie could help open and run the office in Canada / USA since most of our partners and projects are in USA and Canada. London in Canada will most likely be our main North America office, but I find Vancouver and the BC area very startup “friendly”. I believe Natalie and Dena are great leaders and I cannot wait for one of them to be the next CEO. I have great respect for female leaders, starting from my sister, Dena and Natalie, also. In reality, since all of us discuss most decisions, we all have executive partner as job title, since the CEO and other titles, at least for now, are only “official” they do not explain what all of us do. Dena mostly is doing the internal COO / CFO duties (operation and finance), Natalie deals with Partners (CSO / CPO). I work with the technical part of the company (innovation and product / services development). But I do some of Dena’s work in finance related to costs and payments. Dena helps me with the Data Science reporting part and she does a great job on the data visualization part of the data, Dena can do “real magic” with data her visualization. I work with Natalie on the marketing part, I mostly do competitors analyses, what others are doing. Marketing data from our site and other sources. Dena and Natalie both work together when we have a Marketing Event at a conference, Natalie does the “public” part of it, Dena does the “backend” organizational part of the event.

INTERVIEW

Our biggest asset is that the three of us have different opinions about many topics and we complement each other and like to discuss / debate our points until, with brainstorming,, we can come to a compromise. I am first a software engineer, master in business second, Dena is a math geek first, master in business second. Natalie is a communications major first, master in marketing second. I believe our “hybrid” education helped us a lot in “seeing” the same challenge from another angle / vantage point, this helps us a lot in our project work.

If I did not understand business I might have been more concerned with what tech we will use for the new app / service, will it be Microsoft, Red Hat, open source, etc.

As I both a business and engineering major former student, I think about the tech as well, but I also know that the average user does not see what tech I used, he can see only the end result, plus I would have never found out about MVP (most viable product) if I had not studied business, and MVP is better for creating online services, then building a product for months that no one might use later on in the future. With MVP, you get user input fast and that helps build products and services that someone will / might use. We still don't have any public service available for any user, but it will change very fast. We are working on a great AI solution, that would help with everyday problems and management.

[BSD Mag]: Your web page is very simple! The first interpretation that comes to my mind is “sterile”.

[AS]: Yes, the web pages are bad, I agree. It was done years back while the company existed in the business plan. We wanted to give an intro into the data science with a simple and few lines of code solution called ARES; it does a fast search of NASA DB to find info on exoplanets (planets capable of sustaining life). We are still discussing the new page design, we will launch the new stuff when it is ready and right, deadlines are more guidelines for us. We want it perfect. We might launch the new site with our new first public service and the new Data Science blog. The blog will give new users insight into data science, and the public service will give a view to our potential customers and partners into what we can do, cause if a picture says 1000 words, an app says 10000 words.

[BSD Mag]: Why data science? And what does it mean that you are a data scientist?

[AS]: What is data science and a data scientist are very easy questions with complicated answer. Data Scientists are defined as people who are skilled in programming, math and statistical analysis, and data science is the field they work in. I find a simple answer for data science, for me, is using science tools to get useful insight / info from data.

INTERVIEW

I would also say that most data scientists I know are also data natives or have a data native mind set towards the expectations from the apps they use and they create.

For data natives Monica Rogati had a great article on that topic in recode link: <http://www.recode.net/2014/4/10/11625490/the-rise-of-the-data-natives>.

So in short, the digital natives expect the app / tech to work and they program it, data natives expect the app to work as well, but data natives also expect the app to learn on its own so they don't have to program the app. As a data native, I expect my AI app to learn from data it has for me and to make the changes on its own, like what is my favorite song to play when I come at my desk. I expect the AI to learn my daily routines and do them for me, like checking the weather, sending a task summary, checking info on pending projects, check our slack channel chats from my team or company wide chats, remind me about events today, etc. Sending my summary of that info without the need for me to check all of that info on my own. My personal AI is called LEX, LEXA is our Sci Bi AI, used by all. LEX is the bleeding edge version, and LEXA is the stable release version.

[BSD Mag]: What is ARES?

[AS]: We wanted to give an intro into the data science with a simple and few lines of code solution called ARES; it does a fast search of NASA DB to find info on exoplanets (planets capable of sustaining life). ARES was hosted on our webpage, and still is. NASA created R packages that solved the search problems much, much faster, but for the time and given the terabytes of data in the NASA database, ARES was a great solution and great challenge for the time. We might create ARES 2.0, but it will be for some other purpose with NASA's large database, since large data / big data sets are our biggest challenge.

[BSD Mag]: Do you use open source software in your solutions?

[AS]: We are trying to have 100% open source solutions and all will be on git.

The application I use most is Slack and our own build task manager SOPHIE, plus Gitlab and its wonderful management webadmin page and the gitlab slack integration. We do use MEAN JS <http://meanjs.org/>. This is mostly used for our partners or customer's frontend apps; our backend apps that do the data science AI part use R (The R Project for Statistical Computing) <https://www.r-project.org/> with BSD and few times ubuntu servers. I find BSD and its file system is one of the best; since ubuntu started using it in the latest versions, I like what I see. We started using ubuntu server and CoreOs <https://coreos.com/> because of our partners, but Linux compared to BSD had the advantage of having bigger and more active community, at least in the last few years, but BSD, especially PC BSD, will always get my "vote", even compared to any Linux distro. I used GNOME as a main Desktop env, but when I started using Linux Mint, I moved to Cinnamon with PC BSD ver. 10, but I respect what BSD is doing with Lumina desktop.

INTERVIEW

I guess people will say it is still buggy, but it can become a great PC BSD desktop environment in some later version of 10 or version 11.

I use Lumina on the PC BSD Virtual Machine for testing. In the past few months, I have been kinda looking for a good data visualization tool, but it must be open source, so I can change it however I like. For now, good options we use are SVG, d3.js <https://d3js.org/> (under BSD Licence) and Dygraphs <http://dygraphs.com/> charts.JS <http://www.chartjs.org/> (last two under MIT Licence). This web page is a great source for information related to Open Source <https://opensource.com/> and Open Source Solutions and how to use them.

[BSD Mag]: Which open source system is your favorite and why?

[AS]: I can say R on PC-BSD with Firefox Browser since that is what I use daily for 90% of time, but my favorite open source system has to be our own slack add-on task manager named SOPHIE. We named it after my sister and Dena's favorite cousin Sophie. We used slack to do managers for a few years, but after the new year, we moved to SOPHIE. It is still in beta, but we plan to put the entire code or 99% of it on git. It might not be the exact version we are using, but the version will work and anyone can download it and use it and change it for them as they want. We use slack for everything (for most of our internal and partners communication) and SOPHIE is the greatest and most useful tool we had ever build.

My dev team also uses Gitlab CE and the nice integration with slack. I get progress reports on Slack about recent projects without the need to be logged in to gitlab. And SOPHIE puts it all nice and tidy together so I can see the progress the way I want, when I want.

I personally have a MacBookPro 15 2014 model (Apple has superior hardware compared to other vendors for laptops; I have an MBP from 2008 that still works). I used OS X, but I have dual boot with PC BSD. At work, as everyone else, I use System76 Oryx Pro laptops with Linux Mint (<https://www.linuxmint.com/>) (but I have PC BSD as well). Linux Mint is very easy to use, even for the non tech team members, plus it has great support for games; games are great when you have some free time at the office. I personally believe as a team leader that it is not how much time you are at work that counts, it is important how much you work and how much you can get done. I don't want people sitting at the office 9 to 5 doing nothing, it's better to work only 4 hours but really work, and play games for the next 4 hours if you have the time, of course. We don't have an official policy for work hours, except when we have meetings, and everyone is responsible for his / her time. If you don't do the work in the morning, do it in the afternoon or late at night. I usually come to work around 10-11 AM, stay until 21-22, sometimes 2-3 AM, if we have USA CANADA partner meetings, but I will do a nice game break, go to the movies and comeback.

On Saturday evening, we play MMO team games like Star Wars Battlefront at the office.

INTERVIEW

In the past year, I did a lot of traveling across Europe; I spend more than 50% of the time in a different city so I feel like all of Europe is my home, not just London and the UK.

99% of our internal services are browser based, so it really makes no difference what OS we use, except maybe for Microsoft operating systems and using Internet Explorer.

I prefer to use Firefox or Chromium (based on Google Chrome) because of privacy.

[BSD Mag]: You started your career as a software developer. Have you ever been part of an open source community? Why have you decided to go into business direction?

[AS]: In an official capacity or in open source community panels sadly no, for now at least. In high school, I was a math major and did some programming in Pascal and C in my free time. In 2003, I went to the University, and after I got my Bachelor's degree, I enrolled for my first master within the same year, after that within the same year I enrolled into my second master, so from 2003 until 2013, I was working and being a student and being a gamer. From around 2013, we started working on our startup, so full-time schedule from 2003 until now. I helped out within the open source community as user, reporting bugs, helping out with bug reports and answering questions, but I was not part of an open source project team. It might be I still have not found the right project yet, but if I find a good open source project, I will spend time on it maybe within the BSD, CoreOS or Linux Mint community.

If I don't find an interesting open source project to join at the moment, we might start our own open source project since we have a lot of active projects we can put on git, like SOPHIE.

While I was a teaching assistant in Computer Science, I saw the value of a cooperative open source community, but I realized that I wanted more proactive practical projects, not just theoretical. Then I moved into the business world at the telecommunications provider Neotel in the data center, but sadly big companies cannot always use the latest research innovations in the work, even if they use high tech solutions. Neotel started using BSD for the first time for real challenges. I had very wonderful years as a teaching assistant and lecturer at the University and an engineer in Neotel, I learned the value of teamwork in both and both places were great teaching places for me.

Sci Bi is a good combination of both, like the best of both worlds; we use the high best tech solutions and use the newest that science can offer. I "love" my job and enjoy working with Dena, Natalie and the entire Sci Bi team and because of it, I have no problem working 50-70+ hours, since it does not feel like work (at least most of the time).

[BSD Mag]: You have been involved in many projects, you have many certificates, awards and publications. Also, you have graduated from many universities. Why? Do you find it important to gain more diplomas or is it about self-development? Do you think nowadays IT professionals are well educated?

INTERVIEW

Are they interested in gaining knowledge in a different fields and should they be or should their knowledge and experience be focused on one narrow field?

[AS]: Yes, I have worked on many science and research projects (I still do). Most of the certificates are because I found something interesting. I go to the course, and I find that when you know you are going for the certificate, you study more. Most of the certificates were for me not because I need it for some project or work application. The latest Coursera course in 2015 for R programming and data science toolkit, I only took it to see how much that course will help my team to better understand data science, and Coursera and other educational platforms are great, I believe. I personally believe in knowledge sharing and my business degree thesis back in 2009 was investigating the value and potential of web tools in education. It was my belief at the time that online educational platforms that offer blogs, forums and access for everyone will become major educational contributors in the next 10 years. I was wrong, it became popular within 5 years (Coursera one of many).

[BSD Mag]: You like MMO games! Which one have you been playing recently? How do you find time for that?

[AS]: Yes, I believe I mentioned few times during this interview that games are “part” of our Sci Bi company, since we met with Dena during a game, and we found games to help out with stress at the office, and “killing” the down time at the office. We are a small team of about 10 people (most of us work remotely, me being one of them), I don’t know if with 100+ team we can do it the same way but I believe we will always have a game room / break room, maybe in few years VR headsets for everyone, when they become part of every new and modern game. I am a Sci Fi geek and gamer for life and yes I am proud of it.

[BSD Mag]: Are there any challenges your company is facing at the moment? Any plans for the future?

[AS]: Our biggest challenge is to create the perfect publicly available service app web service backed by data science and AI (MVP was a great help here) before September, maybe sooner, and launch the data science blog. It will help us to better and more clearly present our solutions to our partners and the more users we have, the more input we will get from them. SOPHIE’s public version as a Slack add on might do just that.

Yes, our company's biggest and ultimate argument between Dena, Natalie and me: what is better Star Wars (Dena), Star Trek (Natalie), or Stargate (me)? It's our favorite topic for the pub / bar. I can say in this interview that they agreed with me about Stargate being best of Star Wars and Star Trek, but I won't because I will never hear the end of it (plus George being Firefly fan and Lyhn being Battlestar Galactica fan, does not help me a lot, I would say).

INTERVIEW

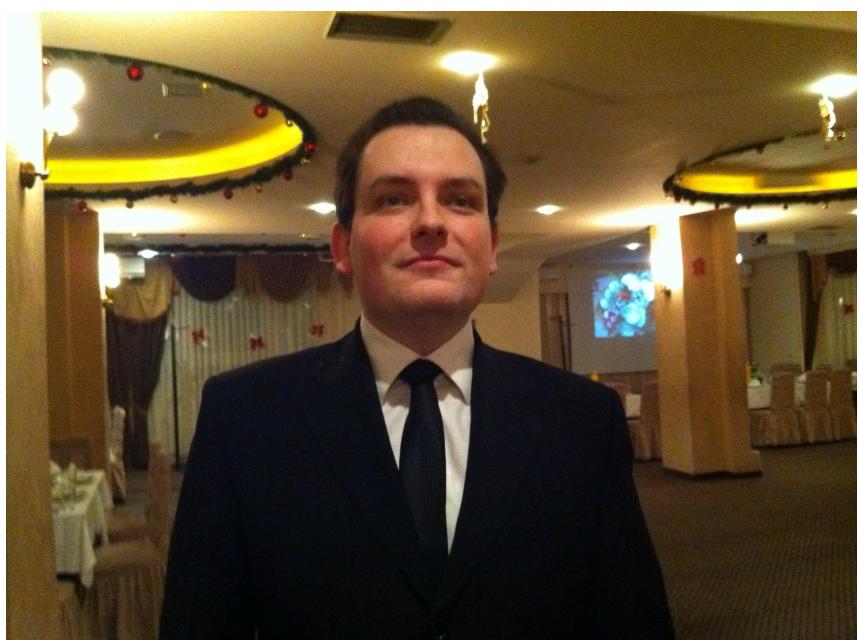
My biggest challenge with my AI engine (LEX and LEXA) is, I am still waiting for an answer on the questions from the hitchhiker's guide to the galaxy question about life and all, and I don't accept 33 as the answer.

[BSD Mag]: Do you have any piece of advice for our readers?

[AS]: My advice is be open minded and try to see the other point of view even if you don't fully agree with it, because finding great solutions requires a team, you cannot do it alone. Any future student that wants to work on great projects, you must learn to be a team player. If you want to be a good leader, you must learn to leave your ego at the door and compromise. I personally believe in leading by example, not by orders or authority. If I want something done, I start to do it first and the others join without the need to order them or boss them around. People can often find me bringing in coffee for everyone at the office like an intern.

My motto in life is: ***Do what you love and you will never work a day in your life.***

About Aleksandar S. Sokolovski:



Aleksandar S. Sokolovski is a data native with years of experience in Software Engineering, Teaching, Lecturing and Data Science. He is a Co-Founder of Sci Bi (startup in offering data science solutions). His biggest asset is being a great mentor and the ability and natural way of knowledge sharing, open source projects are great interpretation of his knowledge sharing philosophy in life. He finds his work in data science and machine learning to be very similar to teaching, since most of it re-

quires training and teaching AI's to use data sets and big data.

He has a great passion for data science and believe that trained AI can help us in our everyday tasks at the office and at home and will give us more quality time with our families.

New web page: <http://www.sci4bi.com>

With the recent British referendum decision to exit from the Europe Union, one of the few opinion polls to correctly predict the outcome was online. What impact is the Internet having on the established political order and is this a force for good?

by Rob Somerville

The UK is currently still in a state of democratic shock after the surprise result to commit to Brexit. Certainly, the general consensus before the result was announced was that it would be a close call, but the majority of people seemed to think that the result would be to remain. What is more surprising, though, is that the traditional polls and surveys generally all agreed with this. While it could be argued that there was no precedent to base a poll on and, therefore, this would undermine sound statistical analysis, there is a much stranger dichotomy at work here. The Internet is generally considered the domain of the younger generation, who pretty much staunchly resided in the remain camp. The older generation (particularly the over 50's), were very much in the Brexit camp. The demographic split was extremely polarised – rich versus poor, North versus South, Globalist versus Isolationist. Where the split was less clear was between the politically left and right; there were advocates for both positions across the political spectrum. So one could logically expect the Internet community to be pro remain, and the poll to reflect this. But the exact opposite was the case.

While the UK tabloids took the position of leave, the broadsheets seemed to take the opposite view. Even the campaign itself – a vicious demonstration of mendacity, dirty tricks, mud slinging and confusion – only addressed the superficial issues. Little focus was truly given in having an intelligent debate about the matter, as both sides rigidly stuck to their respective positions, repeating their superficial mantras until the voting public were thoroughly sick with the whole PR circus. It was only online that the vista became more colourful, from the obvious tin foil hat brigade through to the dogmatists who held their position with an almost religious fervour. The usual online rules applied, of course, from trolling through to ad hominem attacks. While the mood was frequently ugly, and often the boundaries of self control, political correctness and good manners were trampled underfoot in the passionate race to win the argument, at least people were able to share their true observations, opinions and experiences, rather than the sanitised and anodyne discussions that were managed by the traditional PR and communications sectors.

Rob's COLUMN

That is why I believe that the new battlefield for political credibility and votes is online. The politicians have dipped their toes in the water, and while various initiatives have attempted to engage with the Internet audience, as usual, it is too little too late. There are hundreds of thousands of long established political websites, blogs and online videos with millions of visitors and these sites have a unique culture and vision that is rooted in their own particular values, whatever they may be. The audience is committed, loyal and unequivocal in their trust, and in some cases this extends as far as financial support. These voices are trusted as they closely reflect the views and opinions of the target "consumer" far better than the disconnected and out of touch politician. In short, the new kids on the block have not only established a genuine rapport, but have credibility and trust that other forms of media and communications are rapidly losing. Newspaper proprietors and television moguls bemoan the loss of their audience to new media, and like the music industry, refuse to accept the fact that the writing is on the wall. People are fed up with being fed a diet of sterile opinion that refuses to engage with the real issues as the message is passed through the filter of mediocrity. They would rather engage with like minded folk or even tackle the opposite camp head on and have a damn good argument about it, and maybe clarify a few ideas in the process. There is much good that can come from such communities, provided they are not just echo chambers and places to stroke each other's ego. The biggest danger is that we end up with an army of keyboard warriors, that don't get off their bottoms and actually do

something in the real world to change things for the better.

The real battleground however, is trust and credibility. Astroturfing is widespread online, where PR

companies, organisations and individuals set up multiple sites to anonymously promote their point of view, pretending to be grass root activists. And this is the biggest problem online – while the erosion of trust has already undermined traditional politics and the media, the same disease threatens to destroy these groups. Acceptable use policies, moderation of commercial sites and to a small degree, censorship, all have their place, but how can you trust a site or community? Once you emerge from the box of middle ground consensus, how do you deal with divisive and contentious issues such as the EU referendum that have split families and will no doubt end up with a few casualties in the divorce courts?

The biggest problem with new media is that it is far too easy to change your tune at the click of a button. This makes for a very transient and fluid set of ethics in the wrong hands, and while we have tools such as the Wayback Machine, nothing compares to saving pages offline or via a screenshot. A major online UK broadsheet has been caught in the act spiking a negative story on one of the Conservative leadership candidates, and even the Wayback machine suffered from a bad dose of amnesia. If it were not for a few netizens who took the time to document this, no doubt this brazen volte-face would have gone unnoticed.

Rob's COLUMN

The solution to this problem, as ever, is transparency. Especially financial transparency. It is very easy to distil the motives of the opinion maker once you follow the money. Where it becomes more difficult is in the arena of political and social opinion, and those that have great influence as opinion formers. A well presented and engineered video, blog or website can be extremely deceptive. Taken in by quality presentation, powerful argument and rhetoric, it is no wonder that many are taken in by propaganda, and worse, still fall prey to radicalisation. And sadly, while many are abandoning the traditional media hoping to find a more idealistic platform, the risks and potential for the ultimate breakdown of trust are so much greater.